arXiv:1704.04548v1 [cs.LG] 14 Apr 2017

On the Gap Between Strict-Saddles and True Convexity: An $\Omega(\log d)$ Lower Bound for Eigenvector Approximation

Max Simchowitz* Ahmed El Alaoui† Benjamin Recht‡

Abstract

We prove a *query complexity* lower bound on rank-one principal component analysis (PCA). We consider an oracle model where, given a symmetric matrix $M \in \mathbb{R}^{d\times d}$, an algorithm is allowed to make T *exact* queries of the form $w^{(i)} = Mv^{(i)}$ for $i \in \{1,\dots,T\}$, where $v^{(i)}$ is drawn from a distribution which depends arbitrarily on the past queries and measurements $\{v^{(j)}, w^{(j)}\}_{1\le j\le i-1}$. We show that for a small constant ϵ, any adaptive, randomized algorithm which can find a unit vector $\widehat{v}$ for which $\widehat{v}^\top M\widehat{v} \ge (1-\epsilon)\|M\|$, with even small probability, must make $T = \Omega(\log d)$ queries. In addition to settling a widely-held folk conjecture, this bound demonstrates a fundamental gap between convex optimization and "strict-saddle" non-convex optimization [45, 39, 50] of which PCA is a canonical example: in the former, first-order methods can have dimension-free iteration complexity, whereas in PCA, the iteration complexity of gradient-based methods must necessarily grow with the dimension. Our argument proceeds via a reduction to estimating the rank-one spike in a deformed Wigner model. We establish lower bounds for this model by developing a "truncated" analogue of the χ^2 Bayes-risk lower bound of Chen et al. [26].

*UC Berkeley, CA. msimchow@berkeley.edu.
†UC Berkeley, CA. elalaoui@berkeley.edu.
‡UC Berkeley, CA. brecht@berkeley.edu.

1 Introduction

A major open problem in machine learning and optimization is identifying classes of non-convex problems that admit efficient optimization procedures. Motivated by the empirical successes of matrix factorization/completion [86, 68, 49], sparse coding [64], phase retrieval [34] and deep neural networks [48, 12], a growing body of theoretical work has demonstrated that many gradient and local-search heuristics - inspired by convex optimization - enjoy sound theoretical guarantees in a wide variety of non-convex problems [22, 18, 62, 44, 10, 11, 79, 80, 15, 85, 82]. Notably, Ge et al. [39] introduced a polynomial-time noisy gradient algorithm for computing approximate local-minima of non-convex objectives which have the "strict-saddle property": that is, objectives whose first-order stationary points are either local minima, or saddle points at which the Hessian has a strictly negative eigenvalue. It has since been shown that many well-studied non-convex problems can be formulated as "strict saddle" objectives whose local minimizers are all globally optimal (or near-optimal) [78, 50, 40, 15], thereby admitting efficient optimization by local search.

Recently, Jin et al. [45] proposed a gradient algorithm which finds an approximate local minimum of a strict saddle objective in a number of iterations which matches first-order methods for comparable convex problems, up to poly-logarithmic factors in the dimension. This might seem to suggest that, from the perspective of first-order optimization, strict saddle objectives and truly convex problems are identical. But there is a caveat: unlike the algorithm proposed by Jin et al. [45], the iteration complexity of first order methods for optimizing truly convex functions typically *has no explicit dependence on the ambient dimension* [21, 61]. This begs the question:

> Does the iteration complexity of first order methods for strict saddle problems necessarily depend on the ambient dimension? Stated otherwise, is there a gap in the complexity of first-order optimization for "almost-convex" and "truly convex" problems?

This paper answers the above questions in the affirmative by considering perhaps the simplest and most benign strict saddle problem: approximating the top eigenvector of a symmetric matrix $M \in \mathbb{R}^{d\times d}$, also known as rank-one PCA. The latter is best cast as a strict-saddle problem with objective function to be maximized $v \mapsto v^\top M v$, subject to the smooth equality constraint $\|v\|^2 = 1$ [39, 50]. We show that the gradient query complexity of rank-one PCA necessarily scales with the ambient dimension, even in the "easy" regime where the eigengap is bounded away from zero.

More precisely, we consider an oracle model, where given a symmetric matrix $M \in \mathbb{R}^{d\times d}$, an algorithm is allowed to make T *exact* queries of the form $w^{(i)} = Mv^{(i)}$ for $i \in \{1, \ldots, T\}$, where $v^{(i)}$ is drawn from a distribution which may depend arbitrarily on the past queries $\{v^{(j)}, w^{(j)}\}_{1\le j\le i-1}$; these queries are precisely the rescaled gradients of the objective $v \mapsto v^\top M v$. We show (Theorem 2.1) that any adaptive, randomized algorithm which finds a unit vector $\widehat{v}$ for which $\widehat{v}^\top M \widehat{v} \ge \Omega(\|M\|)$ for any symmetric matrix M whose second-eigenvalue is at most γ times its leading eigenvalue in magnitude *must* make at least $T = \Omega(\log d/\log(1/\gamma))$ queries. This matches the performance of the power method and Lanczos algorithms as long as γ is bounded away from one.

In fact, we show that if T is bounded by a small constant times $\log d/\log(1/\gamma)$, then the probability of finding a unit vector $\widehat{v}$ with objective value at least $\Omega(\|M\|)$ is as small as $e^{-d^{\Omega(1)}}$. We also show (Theorem 2.3) that given any $\lambda \ge 2 + \Omega(1)$, it takes $\Omega(\log d/\log\lambda)$ adaptive queries to test if the operator norm of M is above the threshold λ, or below $2 + o(1)$. Our lower bounds are based on the widely studied deformed Wigner random matrix model [51, 33], suggesting that the $\log d$ factor should be regarded as necessary for "typical" symmetric matrices M, not just for some exceptionally adversarial instances.

1.1 Proof Techniques

We reduce the problem of top-eigenvector computation to adaptively estimating the rank-one component θ of the deformation $M = \lambda\theta\theta^\top + W$ of a Wigner matrix W [5, 33, 51] in our query model. Here, λ controls the eigengap, and θ is drawn uniformly from the unit sphere. Unlike many lower bounds for active learning [38, 3, 43], it is insufficient to assume that the algorithm may take the most informative measurements in hindsight, since that would entail estimating θ with only $O(1)$-measurements. As a first pass, we use a recursive application of Fano's method, similar to the strategy adopted in Price and Woodruff [66] for proving lower bounds on adaptive estimation. This method bounds the rate at which information is accumulated by controlling the information obtained from the i-th measurement in terms of the information gained from measurements $1, \ldots, i-1$. Unfortunately, in our setting, this technique can only establish a lower bound of $\Omega(\log d / \log\log d)$ queries.

To sharpen our results, we adopt an argument based on a χ^2-divergence analogue of Fano's inequality, introduced in Chen et al. [26]. But whereas the KL-divergence computations in Fano's inequality allow us to decompose the information obtained at each round into a sum, the adaptivity of the algorithm introduces correlations between the likelihood ratios that appear in the χ^2 computations. Thus, we need to carefully truncate the distributions that arise in our lower bound construction, and restrict them to some carefully-defined "good events". This permit us to bound the rate of information-accumulation.

In general, the probability of these good events conditioned on the spike θ may vary, and thus treating the truncated probabilities as conditional distributions introduces serious complications. To simplify things, we observe that the theory of f-divergences, from which Fano's inequality and the χ^2-analogue in Chen et al. [26] are derived, can be generalized straightforwardly to non-normalized measures, i.e., truncated probability distributions. We therefore derive a general version of the Bayes-risk lower bound from Chen et al. [26] for non-normalized distributions, which, when specialized to χ^2, enables us to prove a sharp lower bound of $\Omega(\log d)$ queries.

To prove the lower bound on testing the spectral norm on $\|M\|$, we reduce the problem to that of testing the null hypothesis $M = W$ for a Wigner matrix W, against an alternative hypothesis $M = \lambda\theta\theta^\top + W$, where θ is drawn uniformly on the sphere and some sufficiently positive λ. The bound mainly follows from Pinsker's inequality (similarly to the combinatorial hypothesis testing lower bound in Addario-Berry et al. [2]) but again with the added nuance of the need to truncate our likelihood ratios due to the adaptivity of the algorithm.

1.2 Related Work

Oracle Lower Bounds for Optimizations. In their seminal work, Nemirovskii and Yudin [60] established lower bounds on the number of calls an algorithm must make to a gradient-oracle in order to approximately optimize a convex function. While they match known upper bounds in terms of dependence on relevant parameters (accuracy, condition number, Lipschitz constant), the constructions are regarded as brittle [8]: the construction considers a worst-case initialization, and makes the strong assumption that the point whose gradient is queried lies in the affine space spanned by the gradients queried up to that iterate. Arjevani and Shamir [9] addresses some of the weakness of the lower bounds [60] (e.g., allowing some randomization), but at the expense of placing more restrictive assumptions of the class of optimization algorithms considered. In contrast, our lower bound places no assumptions on how the algorithm chooses to make its successive queries.

In recent years, lower bounds have been established for stochastic convex optimization [3, 43] where each gradient- or function-value oracle query is corrupted with i.i.d. noise. While these lower bounds are information-theoretic, and thus unconditional, they do not hold in the setting considered in this work, where we are allowed to make exact, noiseless queries. As mentioned above, the proof strategy for proving lower bounds on the exact-oracle model is quite different than in the noisy-oracle setting.

Active Learning and Adaptive Data Analysis. Our proof techniques casts the eigenvector-computation as a type of sequential estimation problem, which have been studied at length in the context of sparse recovery and active adaptive compressed sensing [7, 66, 25, 24]. Due to the noiseless oracle model, our setting is most similar to [66], whereas [7, 25, 24] study measurement noise. Our setting also exhibits similarities to the stochastic linear bandit problem [73]. More broadly, query complexity has received much recent attention in the context of communication-complexity [6, 59], in which lower bounds on query complexity imply corresponding bounds against communication via lifting theorems. Similar ideas also arise in understanding the implication of memory-constraints on statistical learning [77, 76, 70].

Local Search for Non-Convex Optimization. As mentioned in the introduction, there has been a flurry of recent work establishing the efficacy and correctness of local search algorithms in numerous non-convex problems, including Dictionary Learning [11, 80], Matrix Factorization [15, 82, 85], Matrix Completion [40, 16], Phase Retrieval [18, 79, 22], and training neural networks [44]. Particular attention has been devoted to avoiding saddle points in nonconvex landscapes [50, 39, 45, 78], which, without further regularity assumptions, are known to render the task of finding even local minimizers computationally hard [55]. Recent work has also considered second-order algorithms for non-convex optimization [80, 4, 23]. However, to the best of the authors' knowledge, the lower bounds presented in this paper are the first which show a gap in the iteration complexity of first-order methods for convex and "benign" non-convex objectives.

PCA, Low-Rank Matrix Approximation and Norm-Estimation. The growing interest in non-convexity has also spurred new results in eigenvector computation, motivated in part by the striking ressemblence between eigenvector approximation algorithms (e.g. the power method, Lanczos Algorithm [31], Oja's algorithm [71], and newer, variance-reduced stochastic gradient approaches [35, 71]) and analogous first-order convex optimization procedures. Recent works have also studied PCA in the streaming [71], communication-bounded [37, 13], and online learning settings [36]. More generally, eigenvector approximation is widely regarded as a fundamental algorithmic primitive in machine learning [46], numerical linear algebra [31], optimization, and numerous graph-related learning problems [75, 65, 63]. While lower bounds have been established for PCA in the memory- and communication-limited settings [19, 70], we are unaware of lower bounds that pertain to the noiseless query model studied in this work.

Rank-one PCA may be regarded one of the simplest low-rank matrix approximation problems [42]. The numerical linear algebra community has studied low-rank matrix approximation far more broadly, with an eye towards computation-, memory-, and communication-efficient algorithms [58, 57, 67], as well as algorithms which take advantage of the sparsity of their inputs [28, 74, 27]. Previous work has also studied the problem of estimating functions of a matrix's singular values [53], including the special case of estimating Schatten p-norms [52]. To the best of our knowledge, lower bounds for sketching concern the cases where the sketches are chosen non-adaptively.

2 Statement of Main Results

Let $\|\cdot\|$ denote the 2-norm on $\mathbb{R}^d$, and let $\mathcal{S}^{d-1} := \{x \in \mathbb{R}^d : \|x\| = 1\}$ denote the unit sphere. Let $\mathbb{S}^{d\times d}$ denote the set of symmetric $d \times d$ matrices, and for $M \in \mathbb{S}^{d\times d}$, we let $\lambda_1(M) \geq \lambda_2(M) \geq \cdots \geq \lambda_d(M)$ denote its eigenvalues in decreasing order, $v_1(M), v_2(M), \ldots, v_d(M)$ denote the corresponding eigenvectors, and overload $\|M\|$ to denote its operator norm.

Definition 2.1 (Eigenratio)**.** For $\gamma \in [0, 1)$, we define the set of matrices with positive leading eigenvector and bounded *eigenratio* between its first and second eigenvalues:

$$\mathcal{M}_\gamma := \left\{ M \in \mathbb{S}^{d\times d} : \lambda_1(M) = \|M\| > 0, \frac{|\lambda_j(M)|}{\lambda_1(M)} \leq \gamma \quad \forall j \geq 2 \right\}. \tag{2.1}$$

The iteration complexity of rank-one PCA is typically stated in terms of the eigengap $1 - \gamma$. This work instead focuses on lower bounds which hold when the eigengap is close to 1, motivating our parameterization in terms of the eigenratio instead. We now define our query model:

Definition 2.2 (Query Model)**.** And *adaptive query algorithm* Alg with *query complexity* $T \in \mathbb{N}$ is an algorithm which, for rounds $i \in [T]$, queries an oracle with a vector $v^{(i)}$, and receives a noiseless response $w^{(i)} = Mv^{(i)}$. At the end T rounds, the algorithm returns a vector $\widehat{v} = v^{(T+1)} \in \mathcal{S}^{d-1}$. The queries $v^{(i)}$ and output $\widehat{v}$ are allowed to be randomized and adaptive, in that $v^{(i)}$ is a function of $\{(v^{(1)}, w^{(1)}), \ldots, (v^{(i-1)}, w^{(i-1)})\}$, as well as some initial random seed. We say that Alg is *deterministic* if, for all $i \in [T+1]$, $v^{(i)}$ is a deterministic function of $\{(v^{(1)}, w^{(1)}), \ldots, (v^{(i-1)}, w^{(i-1)})\}$. We say that Alg is *non-adaptive* if, for all $i \in [T]$ the distribution of $v^{(i)}$ is independent of the observations $\{(w^{(1)}), \ldots, w^{(i-1)})\}$, (but $\widehat{v}$ may dependent on past observations.)

Example 2.1. The Power Method and Lanczos algorithms [31] are both randomized, adaptive query methods. Even though the iterates $v^{(i)}$ of the Lanczos and power methods converge to the top eigenvector at different rates, they are nearly identical algorithms from our query-complexity perspective: both identify M on the Krylov space $v^{(1)}, Mv^{(1)}, \ldots, M^{T-1}v^{(1)}$. The only difference is that the Lanczos algorithm selects $\widehat{v}$ in a more intelligent manner than the power method. Running the power method from a deterministic initialization would be a *non-randomized* algorithm. Any non-randomized algorithm, even an adaptive one, must take d queries in the worse case, since $d - 1$ queries can only identify a matrix up to a $d - 1$ dimensional subspace. Randomized, but non-adaptive algorithms need to take $\Omega(d)$ queries as well, as established formally in Li et al. [52].

2.1 Lower Bound for Estimation

We let $\mathbb{P}_{\mathsf{Alg}(M)}$ denote probability taken with respect to the randomness of Alg and a fixed M as input, and $\mathbb{P}_{\mathsf{Alg}, M\sim\mathcal{D}}$ denote probability with respect to Alg and M drawn from a distribution $\mathcal{D}$. The main result of this work is the following distributional lower bound:

Theorem 2.1 (Main Theorem)**.** *There exists universal positive constants γ_0, ϵ_0, d_0, c_1 and c_2 such that the following holds: for all $\gamma \leq \gamma_0$, $d \geq d_0$ and $\epsilon \leq \epsilon_0$, there exists a distribution $\mathcal{D}$ supported on $\mathcal{M}_\gamma$ such that the output $\widehat{v}$ of any adaptive query algorithm Alg with query complexity T satisfies*

$$\mathbb{P}_{\mathsf{Alg}, M\sim\mathcal{D}}\left[\langle \widehat{v}, M\widehat{v}\rangle \geq (1-\epsilon)\|M\|\right] \leq c_1 \exp\left\{-c_2 d \cdot (\gamma/\gamma_0)^{2T}\right\}. \tag{2.2}$$

Thus, since a distributional lower bound implies a worst-case lower bound, we have

Corollary 2.2. *Any adaptive T-query algorithm with output $\widehat{v}$ which satisfies*

$$\mathbb{P}_{\mathsf{Alg}(M)}\left[\langle \widehat{v}, M\widehat{v}\rangle \geq \Omega(\|M\|)\right] \geq \Omega(1),$$

for all $M \in \mathcal{M}_\gamma$ must make $T = \Omega(\log d/\log(1/\gamma))$ queries.

The constraint that $\gamma \leq \gamma_0$ implies that there is a large eigengap. In this regime, our lower bound matches the power method, which yields $\widehat{v}$ such that $\widehat{v}^\top M\widehat{v} \geq (1-\frac{1}{\mathrm{poly}(d)})\|M\|$ in $O(\log d/\log(1/\gamma))$ iterations. We also note that, while M is not necessarily positive semi-definite in our construction, one can simply add a multiple of $|\lambda_2(M)| \cdot I$ to enforce this constraint[1], and this only changes γ by a constant.

2.2 Lower Bound for Testing

We now consider the problem of testing whether the operator norm of a symmetric matrix M is below a threshold λ_0, or above a threshold $\lambda_1 > \lambda_0$.

Definition 2.3 (Adaptive Testing Algorithm)**.** An adaptive detection algorithm makes T adaptive, possibly randomized queries as per Definition 2.2, and at the end of T rounds, returns a test $\psi \in \{0,1\}$ which is a function of $\{(v^{(1)}, w^{(1)}), \ldots, (v^{(T)}, w^{(T)})\}$, and some initial random seed.

Our second result establishes a lower bound on the sum of type-I and type-II errors incurred when testing between distributions on matrices with separated operator norms:

Theorem 2.3 (Detection Lower Bound)**.** *There exists universal positive constants λ_0, d_0, c_1, c_2 such that the following holds: for all $d \geq d_0$ and $\lambda \geq \lambda_0$, there exists two distributions $\mathcal{D}_0$ and $\mathcal{D}_1$ on $\mathbb{S}^{d\times d}$ such that*

$$\mathbb{P}_{M\sim\mathcal{D}_0}\left[\|M\| \leq 2 + O\left(d^{-c_1}\right)\right] = 1 \quad \text{and} \quad \mathbb{P}_{M\sim\mathcal{D}_1}\left[\|M\| \geq \lambda - O\left(d^{-c_1}\right)\right] = 1. \tag{2.3}$$

Moreover, for any binary test $\psi \in \{0,1\}$ returned by an adaptive T-query algorithm Alg, we have

$$\mathbb{P}_{\mathsf{Alg},M\sim\mathcal{D}_0}[\psi = 1] + \mathbb{P}_{\mathsf{Alg},M\sim\mathcal{D}_1}[\psi = 0] \geq 1 - \frac{(\lambda/\lambda_0)^T}{d^{c_2}}. \tag{2.4}$$

This implies a worst-case lower bound for testing, matched by the power method for large λ:

Corollary 2.4. *Any randomized adaptive T-query algorithm which can test whether $\|M\| \leq 2+o(1)$ or $\|M\| \geq \lambda = 2 + \Omega(1)$ with probability of error $1/10$ requires at least $T = \Omega(\log d/\log\lambda)$ queries.*

3 Reduction to Estimation

We now construct the distribution used to prove Theorems 2.1 and 2.3. We begin by constructing a family of distributions $\{\mathbb{P}_\theta\}$ on $\mathbb{S}^{d\times d}$, indexed by $\theta \in \mathcal{S}^{d-1}$, and place a prior $\mathcal{P}$ on θ. We then show that if M is drawn from the marginal distribution, then with good probability, M lies in $\mathcal{M}_\gamma$ for an appropriate γ, and that any $\widehat{v}$ for which $\widehat{v}^\top M\widehat{v}$ is large must be close to θ. Hence, establishing the desired lower bound is reduced to a lower bound on estimating θ. The construction is based on the Gaussian Orthogonal Ensemble, also know as the Wigner Model [5].

[1] Proposition 3.2 and the proof of Theorem 3.3 show that $\|\lambda_2(M)\|$ is bounded on $\mathcal{D}$.

Definition 3.1 (Gaussian Orthogonal Ensemble (GOE))**.** We say that $W \sim \mathrm{GOE}(d)$ if the entries $\{W_{i,j}\}_{1 \le i \le j \le n}$ are independent, for $1 \le i < j \le n$, $W_{ij} \sim \mathcal{N}(0,1)$, for $i \in [n]$, $W_{ii} \sim \mathcal{N}(0,2)$, and for $1 \le j < i \le n$, $W_{i,j} = W_{j,i}$. We also define the constant

$$K_d := \frac{1}{\sqrt{d}} \mathbb{E}\left[\|W\|\right] = 2 + d^{-\Omega(1)}. \tag{3.5}$$

For a precise, non-asymptotic upper bound on K_d, we direct the reader to Bandeira and van Handel [14]; an asymptotic bound can be found in Anderson et al. [5], and non-asymptotic bounds with looser constants are shown by Vershynin [83]. We now define the generative process for our lower bound:

Definition 3.2 (Deformed Wigner Model)**.** Let $\lambda > 0$, and $\mathcal{P}$ a distribution supported on $\mathcal{S}^{d-1}$ (e.g., the uniform distribution.) We then independently draw $\theta \sim \mathcal{P}$ and $W \sim \mathrm{GOE}(d)$, and set $M = \lambda \theta\theta^\top + \frac{1}{\sqrt{d}} W$. We also let $\mathbb{P}_u$ denote the law of M conditioned on $\{\theta = u\}$.

In the sequel, we will take our algorithm Alg to be fixed. Abusing notation slightly, we will therefore let $\mathbb{P}_u$ denote the law of M and $\{(v^{(1)}, w^{(1)}), \ldots, (v^{(T)}, w^{(T)}), \widehat{v}\}$ under Alg, conditioned on $\{\theta = u\}$. We now state our main technical result, which establishes a lower bound on estimating θ in the setting of Defintion 3.2, which we prove using Corollary 5.4 in Appendix D.3.

Proposition 3.1 (Main Estimation Result)**.** *Let $\lambda > K_d + o(1)$ and let M be generated from the deformed Wigner model, Definition 3.2. Let $\widehat{v}$ be the output of an adaptive T-query algorithm with input M. Then for any $\eta \ge 0$, we have*

$$\mathbb{E}_{\theta \sim \mathcal{P}} \mathbb{P}_\theta \left[\langle \widehat{v}, \theta \rangle^2 \ge \eta\right] \le \frac{2}{1 - 1/e} \cdot \exp\left\{ \frac{-d\eta}{4\left(c_1 \lambda^2\right)^T} \right\}, \tag{3.6}$$

where $c_1 > 8$ is a universal constant, (observe that $c_1 \lambda^2 > 1$.)

Proposition 3.1 states that, until $\Omega(\log(\eta d)/\log \lambda)$ queries have been made, the probability of having an inner product with θ of at least η is tiny, i.e., $O(e^{-(d\eta)^{\Omega(1)}})$. The following proposition establishes that, if M is drawn from the deformed Wigner model, then with high probability, M lies in $\mathcal{M}_\gamma$ for $\gamma \approx 2/\lambda$, and that optimizing $\langle \widehat{v}, M\widehat{v} \rangle$ entails estimating θ:

Proposition 3.2. *Fix $\delta_0 \in (0,1)$, $\lambda > K_d + 2\sqrt{\log(1/\delta_0)/d}$ and $\theta \in \mathcal{S}^{d-1}$. Let $M \sim \mathbb{P}_\theta$, then the following three assertions simultaneously hold with probability at least $1 - 2\delta_0$:*

1. *$\lambda_{\max}(M) = \|M\| \ge \lambda - 2\sqrt{\log(1/\delta_0)/d}$, and for all $i \ge 2$, $|\lambda_i(M)| \le K_d + 2\sqrt{\log(1/\delta_0)/d}$,*
2. *$M \in \mathcal{M}_\gamma$ for $\gamma = \gamma(d, \lambda, \delta) := \frac{K_d + 2\sqrt{\log(1/\delta_0)/d}}{\lambda - 2\sqrt{\log(1/\delta_0)/d}}$,*
3. *Let $\gamma = \gamma(d, \lambda, \delta)$ as above. For any $\epsilon \le 1 - \gamma$ and any $w \in \mathcal{S}^{d-1}$, if $w^\top M w \ge (1 - \epsilon)\theta^\top M \theta$, then*

$$|\langle w, \theta \rangle| \ge F(\epsilon, \gamma) := \sqrt{\left(\frac{\gamma}{1-\gamma}\right)^2 + 1 - \frac{\epsilon}{1-\gamma}} - \frac{\gamma}{1-\gamma}. \tag{3.7}$$

Remark 3.1. As $\lambda \to \infty$, then $\gamma \to 0$ and $F(\epsilon, \gamma) \to \sqrt{1-\epsilon}$; thus for large λ, optimizing $\langle \widehat{v}, M\widehat{v} \rangle$ is essentially equivalent to estimating θ. The above proposition also lets us take λ to be as small as $K_d + o(1)$, or equivalently, γ arbitrarily close to 1. In this regime, we show in Appendix A.2 that $F(\epsilon, \gamma)$ behaves like $\Theta(1-\gamma)$, provided that $\epsilon = O(1-\gamma)$. Thus, as the eigengap decreases, $\langle \widehat{v}, M\widehat{v} \rangle$ must be ever-closer to $\|M\|$ to ensure that $\widehat{v}$ overlaps with the spike θ. Nevertheless, we can still ensure non-negligible overlap between $\widehat{v}$ and θ for values of γ *arbitrarily close to 1.*

We now state a more detailed version of Theorem 2.1. A formal version of Theorem 2.3 is established in Section 6.

Theorem 3.3 (Formal Statement of Theorem 2.1)**.** *There exist an absolute constant $c_1 > 0$, such that for any $\gamma \in (0, 1/c_1)$ and $\epsilon \in (0, 1-\gamma)$, there exists a distribution $\mathcal{D}$ supported on $\mathcal{M}_\gamma$ such that, for any randomized, adaptive T-query algorithm* Alg*, we have*

$$\mathbb{P}_{M \sim \mathcal{D}, \mathsf{Alg}} \big[\langle \widehat{v}, M\widehat{v} \rangle \geq (1-\epsilon)\|M\| \big] \leq 12 \exp\left\{ -\frac{d}{4} F(\gamma, \epsilon)^2 \cdot (c_1 \gamma)^{2T} \right\}, \tag{3.8}$$

where $F(\gamma, \epsilon)$ is defined in Equation (3.7)*.*

If γ is bounded away from 1 and ϵ bounded away from zero, then by Remark 3.1, the $F(\epsilon, \gamma) = \Omega(1)$, and we recover Theorem 2.1 by observing that the quantity in the exponent is then $-\Omega(d) \cdot (\Omega(\gamma))^{2T}$. However, Theorem 3.3 is more general because, in view of Remark 3.1, it permits γ to be arbitrarily close to 1.

Proof of Theorem 3.3. Set $\lambda = \frac{K_d + 2d^{-1/2}}{\gamma} + 2d^{-1/2} \leq \frac{K_d + 4d^{-1/2}}{\gamma}$. Now we apply Proposition 3.2 with this λ, and with $\delta_0 = 1/e$, so that the conditions of Proposition 3.2 hold with probability at least $1 - 2/e$. If $1 - \langle \widehat{v}, v_1(M) \rangle^2 \leq \epsilon$, then Equation (3.7) implies $\langle \widehat{v}, \theta \rangle^2 \geq \eta := F(\gamma, \epsilon)^2$. Then, we invoke Proposition 3.1 with that value of η, and condition on the event in Theorem 3.2. Finally, we use the bound $\frac{2}{(1-1/e)(1-2/e)} \leq 12$. □

3.1 Computing the Conditional Likelihoods

We begin by introducing some useful simplifications. First, in the spirit of Yao's minimax duality principle [84], we assume that Alg is deterministic[2]. Second, we assume that $v^{(1)}, \ldots, v^{(T)}$ are orthonormal. This is without loss of generality because one can reconstruct the response $w^{(i)}$ to a query $v^{(i)}$ by simply quering the projection of $v^{(i)}$ onto the orthogonal complement of the previous queries $v^{(1)}, \ldots, v^{(i-1)}$, and normalizing. Finally, we introduce a simplification which will make our queries resemble queries of the form $w^{(i)} = \theta \langle v^{(i)}, \theta \rangle +$ i.i.d. noise.

Observation 3.1. For $i \in [T]$ will let P_i then define the orthogonal project onto the complement of the span of $\{v^{(1)}, \ldots, v^{(i)}\}$. We may assume without of generality that, rather that returning responses $w^{(i)} = Mv^{(i)}$, the oracle returns responses $w^{(i)} = P_{i-1} M v^{(i)}$.

[2]Indeed, let $p \in [0,1]$, and E be any event measurable with respect to θ, M, $v^{(1)}, \ldots, v^{(T)}, \widehat{v}$. Then, given a randomized adaptive query algorithm Alg such that $\mathbb{P}_{\theta \sim \mathcal{P}} \mathbb{E}_{\mathsf{Alg}, \mathbb{P}_\theta}[E] \geq p$, we can view Alg as a superposition of deterministic algorithms Alg_ξ, where $\xi \in \Xi$ is a random seed. Then, $\sup_\xi \mathbb{E}_{\theta \sim \mathcal{P}} \mathbb{P}_{\mathsf{Alg}_\xi, M \sim \mathbb{P}_\theta}[E] \geq \mathbb{E}_{\xi \sim \mathcal{D}_\Xi} \mathbb{E}_{\theta \sim \mathcal{P}} \mathbb{P}_{\mathsf{Alg}_\xi, M \sim \mathbb{P}_\theta}[E] = \mathbb{E}_{\theta \sim \mathcal{P}} \mathbb{E}_{\mathsf{Alg}, M \sim \mathbb{P}_\theta}[E] \geq p$.

This is valid because once Alg queries $v^{(1)}, \dots, v^{(i-1)}$, it knows $M(I - P_{i-1})$, and thus, since M and P_{i-1} are symmetric, it also knows $(I - P_{i-1})M$. Thus, throughout, we will take $w^{(i)} = P_{i-1}Mv^{(i)}$. We let also let $Z_i := \{(v^{(i)}, w^{(i)})\}_{1 \le j \le i}$ denote the data collected after the i-th measurement is taken, and let $\mathcal{F}_i$ denote the σ-algebra generated by Z_i. The collection $\{\mathcal{F}_i\}_{1 \le i \le T}$ forms a filtration, and since our algorithm is deterministic, $v^{(i+1)}$ is $\mathcal{F}_i$-measurable. We show that, with our modified measurements $w^{(i)} = P_{i-1}Mv^{(i)}$, then the query-observation pairs $(v^{(i)}, w^{(i)})$ have Gaussian likelihoods conditional on Z_i and u.

Lemma 3.4 (Conditional Likelihoods)**.** *Under $\mathbb{P}_u$, the law of M conditioned on $\theta = u$, we have*

$$P_{i-1}Mv^{(i)} \big| Z_{i-1} \sim \mathcal{N}\left(\lambda(u^\top v^{(i)})P_{i-1}u, \frac{1}{d}\Sigma_i\right), \quad \text{where} \quad \Sigma_i := P_{i-1}\left(I_d + v^{(i)}v^{(i)\top}\right)P_{i-1}. \tag{3.9}$$

In particular, $w^{(i)}$ is conditionally independent of $w^{(1)}, \dots, w^{(i-1)}$ given Z_{i-1} and $\theta = u$.

Lemma 3.4 is proved in Appendix B. We remark that Σ_i is rank-deficient, with its kernel being equal to the span of $\{v^{(1)}, \dots, v^{(i-1)}\}$. Nevertheless, because the mean vector $\lambda(u^\top v^{(i)})P_{i-1}u$ lies in the orthogonal complement of $\ker \Sigma_i$, computing $\Sigma_i^{-1}(\lambda(u^\top v^{(i)})P_{i-1}u)$ can be understood as $\Sigma_i^\dagger(\lambda(u^\top v^{(i)})P_{i-1}u)$, where $\dagger$ denotes the Moore-Penrose pseudo-inverse [42]. We write

$$\|v\|_\Sigma^2 := v^\top \Sigma^\dagger v. \tag{3.10}$$

and we will use the following equality and inequality frequently and without comment:

$$\|\lambda v\|_{\frac{1}{d}\Sigma_i}^2 = d\lambda^2 \|v\|_{\Sigma_i}^2, \quad \text{and} \quad \|P_{i-1}u\|_{\Sigma_i}^2 \le \|u\|^2. \tag{3.11}$$

These just follow from the facts that P_{i-1} is an orthogonal projection and $\Sigma_i \succeq P_{i-1}$, and so $P_{i-1}\Sigma_i^\dagger P_{i-1} \preceq I$.

4 A First Attempt: a Lower Bound of $\Omega(\log d / \log\log d)$

Many adaptive estimation lower bounds are often shown by considering the most informative measurements an algorithm could take if it knew the true hidden parameter [38, 43, 3, 72]. Unfortunately, this line of attack in insufficient for a non-vacuous lower bound in our setting: if an oracle tells the algorithm to measure at a unit vector v for which $\langle v, \theta \rangle$ is at least $\Omega(1)$, then we would have $\langle v, Mv \rangle = \Omega(\|M\|)$, and so by Proposition 3.2, we would verify that v is close to θ. Of course, it is highly unlikely that our *first* measurement $v^{(1)}$ is close to the true θ; indeed, if θ is drawn uniformly from the sphere $\mathcal{S}^{d-1}$, then $\langle v, \theta \rangle^2 = O(1/d)$ with high probability. But what about the *second* measurement, or the *third*? What is to stop the algorithm from rapidly learning to take highly informative measurement? To show this cannot happen, we will adopt a simple recursive strategy:

1. We relate the information collected at stage k to the inner products $\langle v_i, \theta \rangle^2$, $i \in [k]$.
2. We bound the inner product of θ with the $k+1$-st query by its inner products with all past queries as

$$\langle v^{(k+1)}, \theta \rangle^2 \le C(\lambda) \cdot \sum_{i=1}^{k} \langle v^{(i)}, \theta \rangle^2,$$

where $C(\lambda)$ is a constant depending on λ.

To demonstrate the above proof strategy, we start by establishing a sub-optimal lower bound of $\Omega(\log d/\log\log d)$. Then in Section 5, we introduce a more refined machinery to sharpen the bound to $\Omega(\log d)$. First, we observe that the mutual information $I(Z_k;\theta)$ between θ and Z_k (see, e.g., Cover and Thomas [29]) is controlled by the inner products $\langle v^{(i)},\theta\rangle^2$, $i\in[k]$:

Proposition 4.1. *Let $\mathcal{P}$ be an isotropic probability distribution supported on $\mathcal{S}^{d-1}$, and $\mathbb{P}_u$ denote the law of $M=\lambda\theta\theta^\top+\frac{1}{\sqrt{d}}W$ conditioned on $\theta=u$. Then for all integers $k\geq 1$,*

$$I(Z_k;\theta)\leq\frac{\lambda^2}{2}\left(k+\sum_{i=1}^k\tau_i\right),\quad \textit{where}\quad \tau_i:=\mathbb{E}_{u\sim\mathcal{P}}\,\mathbb{E}_{\mathbb{P}_u}\left[d\cdot\langle v^{(i)},u\rangle^2\right]. \tag{4.12}$$

We prove Proposition 4.1 in Appendix C.1. We now recursively bound the mutual information $I(Z_k;\theta)$, using an argument similar to Price and Woodruff [66], with the exception that we will rely on a more recent continuum formulation of Fano's inequality [32, 26] to control the information :

Proposition 4.2 (Global Fano [32])**.** *Let $\mathcal{P}$ be a prior over a measure space $(\Theta,\mathcal{G})$, and let $\{\mathbf{P}_\theta\}$ denote a family of distributions over a space $(\mathcal{X},\mathcal{F})$ indexed by $\theta\in\Theta$. Then, if $\mathcal{A}$ is an action space, $\mathcal{L}:\mathcal{A}\times\mathcal{X}\to\{0,1\}$ is a loss function, and $\mathfrak{a}:\mathcal{X}\mapsto\mathcal{A}$ is a measurable map, we have*

$$\mathbb{E}_{\theta\sim\mathcal{P}}\mathbb{P}_{X\sim\mathbf{P}_\theta}\left[\mathcal{L}(\mathfrak{a}(X),\theta)=0\right]\leq\frac{I(X;\theta)+\log 2}{\log\left(1/\sup_{a\in\mathcal{A}}\mathbb{P}_{\theta\sim\mathcal{P}}[\{L(a,\theta)=0\}]\right)}. \tag{4.13}$$

We will apply Proposition 4.2 at each query stage k: we let $\Theta=\mathcal{S}^{d-1}$, θ denote the rank-one spike, $\mathcal{P}$ the prior over θ, X to be the data $Z_k=\{(v^{(1)},w^{(1)}),\ldots,(v^{(k)},w^{(k)})\}$ collected at the end of round k, and $\mathbf{P}_u$ to be the law of $X=Z_k$ conditioned on $\theta=u$. We use the action space $\mathcal{A}=\mathcal{S}^{d-1}$, our actions will be the $k+1$-st query, $v^{(k+1)}$, and the loss function we consider is

$$\mathcal{L}(v^{(k+1)},\theta)=\mathbb{I}\left(\langle v^{(k+1)},\theta\rangle^2\geq\tau\right), \tag{4.14}$$

for some fixed τ. This leads to the following bound, proved in Appendix C.1.

Proposition 4.3. *Let $\mathcal{P}$ denote an isotropic distribution on the sphere $\mathcal{S}^{d-1}$, which satisfies the following concentration bound for some constants $C_1,C_2>0$ and all $t\geq 0$,*

$$\sup_{v\in\mathcal{S}^{d-1}}\mathbb{P}_{u\sim\mathcal{P}}\left[d\cdot\langle v,u\rangle^2\geq t\right]\leq e^{C_2-C_1t}. \tag{4.15}$$

Then, the sequence $\left\{\tau_i:=\mathbb{E}_{u_0\sim\mathcal{P}}\mathbb{E}_{\mathbb{P}_{u_0}}\left[d\cdot\langle v^{(i)},u_0\rangle^2\right],\ i\geq 1\right\}$ satisfies the following recursion: for all $k\geq 1$, $t>C_2/C_1$, we have

$$\mathbb{E}_{u_0\sim\mathcal{P}}\mathbb{P}_{u_0}\left[d\cdot\left\langle v^{(k+1)},u_0\right\rangle^2\geq t\right]\leq\mathbb{I}(t\leq d)\cdot\left(1\wedge\frac{\log 2+\frac{\lambda^2}{2}(k+\sum_{i=1}^k\tau_i)}{C_1t-C_2}\right). \tag{4.16}$$

Therefore, integrating over t yields

$$\tau_{k+1}\ \leq\ \frac{C_2}{C_1}+\frac{\log 2+\frac{\lambda^2}{2}(k+\sum_{i=1}^k\tau_i)}{C_1}\left(1+\log\frac{C_1d}{\log 2+\frac{\lambda^2}{2}(k+\sum_{i=1}^k\tau_i)}\right). \tag{4.17}$$

In Appendix F, we prove that if $\mathcal{P}$ is the uniform distribution on $\mathcal{S}^{d-1}$, then we can take $C_1 = 1/8$, and $C_2 = 4$ in the above proposition. This relies on the following concentration result:

Lemma 4.4 (Spherical Concentration)**.** *Let $\theta \sim \mathcal{P}$ where $\mathcal{P}$ is the uniform distribution over $\mathcal{S}^{d-1}$. Then for all $v \in \mathcal{S}^{d-1}$,*

$$\mathbb{P}_{\theta\sim\mathcal{P}}\left[\sqrt{d}\cdot|\langle v,\theta\rangle| \geq \sqrt{2}+t\right] \leq e^{-t^2/2}. \tag{4.18}$$

Hence, from Equation (4.17), τ_k grows by at most $O(\log d)$ after each query. Hence, until $\log d/\log\log d$ queries are taken, τ_k will be $o(d)$, which entails that θ will not be accurately estimated. It is worth understanding why this spurious $\log\log d$ factor appears using KL. The main weakness with Proposition 4.2 is that the denominator contains the logarithm of the "best-guess probability" $\log\left(\sup_{a\in\mathcal{A}}\mathbb{P}_{\theta\sim\mathcal{P}}[\{L(a,\theta)=0\}]\right)$. This results in a very weak tail bound on $\mathbb{P}(\tau_k > t)$ of $O(1/t)$, which incurs a $\log d$ factor when integrated. To overcome this weakness, we will work instead with estimates based on the χ^2 divergence, which will be a lot more careful in taking advantage of the small value of $\sup_{a\in\mathcal{A}}\mathbb{P}_{\theta\sim\mathcal{P}}[\{L(a,\theta)=0\}]$.

5 A Sharper Lower Bound on Estimation

In this section, we use more refined machinery based on the χ^2-divergence to sharpen the lower bound from $\Omega(\log d/\log\log d)$ to $\Omega(\log d)$. When bounding the KL-divergences in Proposition 4.1, the proof crucially relies upon the fact that the log-likelihoods decompose into a sum, and could thus be bounded using linearity of expectations. This is no longer the case when working with the squares of likelihood-ratios which arise in the χ^2 divergence, because the adaptivity of the queries can introduce strong correlations between likelihood ratios arising from subsequent measurements. To remedy this, we will proceed by designing a sequence of "good truncation events" for each $\theta \in \mathcal{S}^{d-1}$, on which the likelihood ratios will be well-behaved. We now fix some positive numbers $\tau_1, \ldots, \tau_{T+1}$ to be specified later, and define the events A_u^k for $u \in \mathcal{S}^{d-1}$ and integer k by

$$A_u^k = \left\{\forall i \in [k] : d\cdot\langle u, v^{(i)}\rangle^2 \leq \tau_i\right\}. \tag{5.19}$$

For an arbitrary probability measure $\mathbb{P}$ on a space $(\mathcal{X},\mathcal{F})$, and an event $A \in \mathcal{F}$, we use the following notation to denote the truncated (non-normalized) measure

$$\mathbb{P}\left[B; A\right] := \mathbb{P}\left[B\cap A\right], \qquad \forall B \in \mathcal{F}. \tag{5.20}$$

In the sequel, we will be working with the measures $\mathbb{P}_u[;A_u^k]$. Note that these are no longer actual probability measures, since their total mass is $\mathbb{P}_u[A_u^k]$, which is in general strictly less than one. In Appendix E, we show that f-divergences - a family of measures of distance between distributions which include both the KL and the χ^2 divergence [30, 26] - generalize straightforwardly to non-negative measures which are not normalized (e.g., truncated probability distributions.) Leaving the full generality to the appendix, we will use a "generalized f-divergence" between non-normalized measures which modifies the classical χ^2-divergence (for a comparison to the classical χ^2 divergence, see the discussion following Remark E.1.)

Definition 5.1 (χ^2+1-divergence)**.** Let $\mathbb{P}, \mathbb{Q}$ denote two nonnegative measures on a space $(\mathcal{X}, \mathcal{F})$, such that $\mathbb{Q}[\mathcal{X}] > 0$, and $\mathbb{P}$ is absolutely continuous with respect to $\mathbb{Q}$[3]. We define

$$D_{\chi^2+1}(\mathbb{P}, \mathbb{Q}) := \int \left(\frac{\mathrm{d}\mathbb{P}}{\mathrm{d}\mathbb{Q}}\right)^2 \mathrm{d}\mathbb{Q}. \tag{5.21}$$

When $\mathbb{Q}$ is a probability distribution, the above can be written as $\mathbb{E}_{\mathbb{Q}}\left[\left(\frac{\mathrm{d}\mathbb{P}}{\mathrm{d}\mathbb{Q}}\right)^2\right]$.

In Appendix E, we prove a generalization of the f-divergence Bayes risk lower bounds of Chen et al. [26], which we state here for the χ^2+1 divergence:

Proposition 5.1. *Adopting the setup of Proposition 4.2, let $\mathcal{P}$ be a distribution over a space Θ, $\{\mathbf{P}_\theta\}_{\theta\in\Theta}$ be a family of probability measures on $(\mathcal{X}, \mathcal{F})$, $\mathcal{A}$ be an action space, $\mathcal{L} : \mathcal{A} \times S^{d-1} \to \{0, 1\}$ a binary loss function, and let $\mathfrak{a}$ denote a measurable map from $\mathcal{X}$ to $\mathcal{A}$. Given a family $\{A_\theta\}_{\theta\in\Theta}$ of $\mathcal{F}$-measurable events, let $\mathbf{P}_\theta[\cdot; A_\theta]$ denote the truncated measure as per Equation* (5.20)*. Set*

$$V^{\mathfrak{a}} := \mathbb{E}_{\theta\sim\mathcal{P}}\mathbf{P}_\theta\left[\{\mathcal{L}(\mathfrak{a}(X), \theta) = 0\}; A_\theta\right], \qquad \text{and} \quad V_0 := \sup_{a\in\mathcal{A}} \mathbb{P}_{\theta\sim\mathcal{P}}\left[\{\mathcal{L}(a, \theta) = 0\}\right].$$

Then, for any nonnegative measure $\mathbf{Q}$ on $(\mathcal{X}, \mathcal{F})$, we have

$$V^{\mathfrak{a}} \le V_0 + \sqrt{V_0(1 - V_0)\mathbb{E}_{\theta\sim\mathcal{P}} D_{\chi^2+1}\left(\mathbf{P}_\theta\left[\cdot; A_\theta\right], \mathbf{Q}\right)}. \tag{5.22}$$

Remark 5.1. Even though the above proposition is an analogue of Corollary 7 in Chen et al. [26], it cannot be derived merely as a consequence of that bound, and is sharper and easier to use than bounds that would arise by replacing the truncated distributions $\mathbf{P}_\theta[\cdot; A_\theta^k]$ with conditional distributions $\mathbf{P}_\theta[\cdot\big|A_\theta^k]$. See Remark E.2 for further discussion.

As in Section 4, we take $\Theta = \mathcal{S}^{d-1}$, $\mathbf{P}_u$ to be the distribution of Z_k conditioned on $\{\theta = u\}$, and as in our above discussion, A_θ^k will define the truncation events from Equation (5.19); the index k for which we apply Proposition 5.1 will always be clear from context. If we were to follow the KL case, we would bound the corresponding mutual information quantity $\mathbb{E}_{\theta\sim\mathcal{P}} D_{\chi^2+1}(\mathbf{P}_\theta[\cdot; A_\theta], \mathbf{Q})$ by taking $\mathbf{Q}$ to be the law of Z_k induced by M when $M \sim \mathbb{E}_{\theta\sim\mathcal{P}}\mathbb{P}_\theta$. We would then apply Jensen's inequality (in view of Lemma E.1) to upper bound the corresponding mutual information quantity by the average χ^2 + 1-divergence $\mathbb{E}_{u_0,u_1\sim\mathcal{P}} D_{\chi^2+1}(\mathbf{P}_{u_0}[\cdot; A_{u_0}], \mathbf{P}_{u_1}[\cdot; A_{u_1}])$ between $\mathbf{P}_{u_0}$ and $\mathbf{P}_{u_1}$, where u_0 and u_1 are both drawn i.i.d. from $\mathcal{P}$.

This argument does not work in our setting, because once we restrict to the events A_θ^k, two measures $\mathbf{P}_{u_0}[\cdot; A_{u_0}^k]$ and $\mathbf{P}_{u_1}[\cdot; A_{u_1}^k]$ may no longer be absolutely continuous, and thus have an infinite χ^2+1-divergence. Instead, we apply Proposition 5.1 with the measure $\mathbf{Q} := \mathbf{P}_0$ to denote the (un-truncated) probability law of Z_k under the random matrix $M = \frac{1}{\sqrt{d}}W$. Since $\mathbf{P}_0$ is un-truncated, and since the GOE matrix W has a continuous density, all the measures $\mathbf{P}_u$, and thus $\mathbf{P}_u[\cdot; A_u^k]$, are absolutely continuous with respect to it. Thus, we can use the events A_u^k to control the χ^2+1 divergence as follows:

[3]That is, for every $A \in \mathcal{F}$, $\mathbb{Q}[A] = 0$ implies that $\mathbb{P}[A] = 0$

Lemma 5.2 (Upper Bound on Likelihood Ratios)**.** *Let $u \in \mathcal{S}^{d-1}$, and A_u^k be the event in Equation* (5.19)*. For $i \geq 1$ and $v^{(1)}, \cdots, v^{(i)} \in \mathcal{S}^{d-1}$, define the expected conditional likelihood ratio*

$$g_i\left(u; \{v^{(j)}\}_{1\leq j\leq i}\right) \;:=\; \mathbb{E}_{\mathbf{P}_0}\left[\left(\frac{\mathrm{d}\mathbf{P}_u(Z_i|Z_{i-1})}{\mathrm{d}\mathbf{P}_0(Z_i|Z_{i-1})}\right)^2 \Big| \{v^{(j)}\}_{1\leq j\leq i}\right]. \tag{5.23}$$

Moreover, let $\mathcal{V}_\theta^k := \left\{(v^{(1)}, \ldots, v^{(k)}) \in (\mathcal{S}^{d-1})^k \;:\; \forall i \in [k],\; d\cdot\langle v^{(i)}, \theta\rangle^2 \leq \tau_i\right\}$. Then,

$$D_{\chi^2+1}(\mathbf{P}_u, \mathbf{P}_0) = \mathbb{E}_{\mathbf{P}_0}\left[\left(\frac{\mathrm{d}\mathbf{P}_u(Z_k; A_u^k)}{\mathrm{d}\mathbf{P}_0(Z_k)}\right)^2\right] \;\leq\; \sup_{v^{(1)},\ldots,v^{(k)}\in\mathcal{V}_u^k} \prod_{i=1}^k g_i\left(u; \{v^{(j)}\}_{1\leq j\leq i}\right). \tag{5.24}$$

The above Lemma specializes Lemma C.2, proved in Appendix C.2. Noting that the conditional laws $\mathbf{P}_u(Z_i|Z_{i-1})$ and $\mathbf{P}_0(Z_i|Z_{i-1})$ have Gaussian densities, a computation detailed in Lemma C.3 yields

$$g_i\left(u; \{v^{(j)}\}_{1\leq j\leq i}\right) \;=\; \exp\left\{\lambda^2 \cdot d(u^\top v^{(i)})^2 \cdot \|P_i u\|_{\Sigma_i}^2\right\}. \tag{5.25}$$

Using the bound $\|P_i u\|_{\Sigma_i}^2 \leq \|u\|^2 = 1$ (Equation (3.11)), and that $d\left(u^\top v^{(i)}\right)^2 \leq \tau_i$ on A_u^k, we get

$$\mathbb{E}_{\mathbf{P}_0}\left[\left(\frac{\mathrm{d}\mathbf{P}_u(Z_k; A_u^k)}{\mathrm{d}\mathbf{P}_0(Z_k)}\right)^2\right] \;\leq\; \sup_{v^{(1)},\ldots,v^{(k)}\in\mathcal{V}_u^k} \prod_{i=1}^k g_i\left(u; \{v^{(j)}\}_{1\leq j\leq i}\right) \leq e^{\lambda^2\sum_{i=1}^k \tau_i}. \tag{5.26}$$

Combining these bounds allows us to prove Theorem 5.3 below, which establishes an analogue of the recursion given by Proposition 4.3.

Theorem 5.3. *Let $\mathbb{P}_\theta$ denote the law of the sequential query model from the matrix $M = \lambda\theta\theta^\top + \frac{1}{\sqrt{d}}W$, where $W \sim \mathrm{GOE}(d)$, and $\theta \sim \mathcal{P}$, where $\mathcal{P}$ is a distribution supported on $\mathcal{S}^{d-1}$. Let $v^{(1)}, \ldots, v^{(T)}, \widehat{v}$ denote the queries and output of a deterministic algorithm* Alg*. Then for any $\tau_1, \ldots, \tau_{T+1} > 0$,*

$$\begin{aligned}\mathbb{E}_{\theta\sim\mathcal{P}}\mathbb{P}_\theta\left[\exists k \in [T+1] : d\langle v^{(k)}, \theta\rangle^2 > \tau_k\right] \leq \sup_{v\in\mathcal{S}^{d-1}} \mathbb{P}_{\theta\sim\mathcal{P}}\left[d\langle v, \theta\rangle^2 \geq \tau_1\right] \\ + 2\sum_{k=2}^{T+1} e^{\frac{\lambda^2}{2}\sum_{i=1}^{k-1}\tau_i} \cdot \sqrt{\sup_{v\in\mathcal{S}^{d-1}} \mathbb{P}_{\theta\sim\mathcal{P}}\left[d\langle v, \theta\rangle^2 \geq \tau_k\right]}.\end{aligned} \tag{5.27}$$

The above theorem gives an upper bound on the probability that the inner products $\langle v^{(i)}, \theta\rangle^2$ grow faster than the sequence $\{\tau_i\}$. The summands on the right-hand side of Equation (5.27) exhibit a tradeoff between the values $\tau_1, \ldots, \tau_{k-1}$, which capture the information gathered up to round k, and, as in Section 4, the term $\sup_{v\in\mathcal{S}^{d-1}} \mathbb{P}_{\theta\sim\mathcal{P}}[d\langle v, \theta\rangle^2 \geq \tau_k]$, which represent how likely the best guess of θ one could make without making any queries. We prove this theorem in Appendix D.1

When we specialize Theorem 5.3 by taking $\mathcal{P}$ to be the uniform measure on the sphere, we obtain the following corollary, established in Appendix D.2.

Corollary 5.4. *Let $\mathcal{P}$ be the uniform measure on the sphere $\mathcal{S}^{d-1}$. Then, for any $\delta \in (0,1)$,*

$$\mathbb{E}_{\theta\sim\mathcal{P}}\mathbb{P}_\theta\left[\exists k \in [T+1] : \langle v^{(k)},\theta\rangle^2 \geq \left(2\lambda^2 \cdot c(\delta,\lambda)\right)^{k-1} \cdot \frac{2\left(\sqrt{\log(1/\delta)}+1\right)^2}{d}\right] \quad \leq \quad \frac{2\delta}{1-\delta}, \quad (5.28)$$

where $c(\delta,\lambda) = (1+1/\lambda^2)\left\{(1-1/2\lambda^2)(1-\sqrt{1/(1+\log(1/\delta))})\right\}^{-1} = 1 + o(1/\lambda) + o(\delta)$.

We observe that in Corollary 5.4, the inner products $\langle v^{(k)},\theta\rangle^2$ are unlikely to grow faster than $\Omega\left(\frac{\lambda^{2(k-1)}}{d}\right)$ with high probability. Thus, in order for $\langle v^{(k)},\theta\rangle^2$ to be at least $\Omega(1)$ at one of the T iterates, one needs at least $T = \Omega(\log d/\log\lambda)$ iterations. This is the insight which underlies the proof of the main technical result, Proposition 3.1, which is deferred to Appendix D.3.

Examining the proof of Corollary 5.4, the base of $2\lambda^2$ in the exponent is essentially the best we can hope from our techniques. While this base leads to the order optimal $\log d/\log\lambda$ sample complexity in the large λ regime, we see that as λ approaches $K_d \approx 2$, and thus our upper bound on γ from Proposition 3.2 approaches 1, our lower bound still permits $\langle v^{(i)},\theta\rangle^2$ to grow at a rate of, say, 8^k. A classical result in random matrix theory [33] states that the first eigenvector of M correlates with the spike θ as soon as $\lambda > 1$. In this regime, our lower bound allows $\langle v^{(i)},\theta\rangle^2$ to grow at a rate of about 2^T. This appears to be loose, because it a) a does not rule out a fast rate of convergence despite a vanishing eigengap and b) does not reflect that the rank one perturbation $\theta\theta^\top$ no longer correlates with θ if $\lambda < 1$ [33].

6 Testing the Operator Norm

We retain the notation of the previous section, letting A_u^k denote the truncation events from Equation (5.19), and $\mathbb{P}_u[\cdot; A_u^k]$ denote the corresponding truncated measures. We are interested in testing whether M is drawn from $\mathbb{P}_0$, or from $\mathbb{P}_\theta$ for some θ and $\lambda > 0$. Because the test is a measurable function of Z_T, we establish a lower bound on testing between the distribution $\mathbf{P}_0$ and $\overline{\mathbf{P}} := \mathbb{E}_{\theta\sim\mathcal{P}}\mathbf{P}_\theta$; where again, for $\theta \in \{0\}\cup\mathcal{S}^{d-1}$, $\mathbf{P}_\theta$ denotes the law of Z_T induced by $M = \lambda\theta\theta^\top + \frac{1}{\sqrt{d}}W$. Because this test only requires one bit of information, we need to show that TV-distance between $\mathbf{P}_0$ and $\overline{\mathbf{P}}$ is $o(1)$ until sufficiently many queries have been made. It will therefore be insufficient to bound the TV-distance between $\overline{\mathbf{P}}$ and $\mathbf{P}_0$ by the average TV-distance between $\mathbf{P}_\theta$ and $\mathbf{P}_0$ as in the last section, since testing between the two cases with good probability queries only $O(1)$ queries. Hence, we will more carefully bound TV between $\overline{\mathbf{P}}$ and $\mathbf{P}_0$ by applying Pinsker's inequality [81] to the null and mixture distributions. As in the previous section, we will also need to use truncation.

Proposition 6.1 (Truncated χ^2 Inequality)**.** *Let $\mathcal{P}$ be a distribution over a space Θ, $\{\mathbf{P}_\theta\}_{\theta\in\Theta}$ be a family of probability measures on $(\mathcal{X},\mathcal{F})$. Define the marginal distribution $\overline{\mathbf{P}}$ on $(\mathcal{X},\mathcal{F})$ and its restriction $\overline{\mathbf{P}}[B; A_\theta]$ via*

$$\overline{\mathbf{P}}[B] = \mathbb{E}_{\theta\sim\mathcal{P}}\mathbf{P}_\theta[B], \quad \textit{and} \quad \overline{\mathbf{P}}\left[B; A_\theta^k\right] = \mathbb{E}_{\theta\sim\mathcal{P}}\mathbf{P}_\theta\left[B\cap A_\theta^k\right]. \quad (6.29)$$

Then, if $p = \overline{\mathbf{P}}[\mathcal{X}; A_\theta^k]$, we have for any probability measure $\mathbf{Q}$ on $(\mathcal{X},\mathcal{F})$

$$\|\mathbf{Q}-\overline{\mathbf{P}}\|_{\mathrm{TV}} \leq \frac{1}{2}\sqrt{\mathbb{E}_{\mathbf{Q}}\left[\left(\frac{\mathrm{d}\overline{\mathbf{P}}[\cdot; A_\theta^k]}{\mathrm{d}\mathbf{Q}}\right)^2\right] - 1} + \frac{\sqrt{2(1-p)}+(1-p)}{2}. \quad (6.30)$$

Remark 6.1. As in Proposition 5.1, dispensing with the normalization constants by considering truncated rather than conditional distributions greatly simplifies the presentation. Unlike Proposition 5.1, the Proposition 6.1 is established elementarily, without f-divergence machinery.

As alluded to above, we will take $\mathbf{Q} = \mathbf{P}_0$, and $\mathbf{P}_\theta$ to be the law of Z_T under $M = \lambda\theta\theta^\top + \frac{1}{\sqrt{d}}W$. We take $A_\theta = A_\theta^T$ defined in Equation (5.20). By definition of $\overline{\mathbf{P}}[\cdot; \{A_\theta\}]$ and Fubini's theorem, we write

$$\mathbb{E}_{\mathbf{Q}}\left[\left(\frac{\mathrm{d}\overline{\mathbf{P}}[\cdot; A_\theta]}{\mathrm{d}\mathbf{Q}}\right)^2\right] = \mathbb{E}_{\theta,\theta'\sim\mathcal{P}}\mathbb{E}_{\mathbf{P}_0}\left[\frac{\mathrm{d}\mathbf{P}_\theta[\cdot; A_\theta^T]\mathrm{d}\mathbf{P}_{\theta'}[\cdot; A_{\theta'}^T]}{(\mathrm{d}\mathbf{P}_0)^2}\right]. \tag{6.31}$$

Equation (6.31) is a standard observation in combinatioral hypothesis testing [2], and known as the (conditional) second moment method in probabilistic combinatorics [1]. The quantity on the right hand side of Equation (6.31) resembles the $\chi^2 + 1$ divergence, but with a product of two different likelihoods in the numerator. Because these likelihoods take large values on different parts of the space, Equation (6.31) is typically much smaller than the $\chi^2 + 1$ divergence between $\mathbf{P}_\theta$ and $\mathbf{P}_0$. We now specialize Lemma C.2 in the appendix to establish a bound on Equation (6.31).

Lemma 6.2 (Upper Bound on Likelihood Ratios)**.** *Given $u_1, u_2 \in \mathcal{S}^{d-1}$, and let $A_{u_1}^k$ and $A_{u_2}^k$ be events as in Equation* (5.19)*. For $i \geq 1$ and $v^{(1)}, \cdots, v^{(i)} \in \mathcal{S}^{d-1}$, define the expected conditional likelihoods ratios*

$$g_i\left(u_1, u_2; \{v^{(j)}\}_{1\leq j\leq i}\right) \;:=\; \mathbb{E}_{\mathbf{P}_0}\left[\frac{\mathrm{d}\mathbf{P}_{u_1}(Z_i|Z_{i-1})\mathrm{d}\mathbf{P}_{u_2}(Z_i|Z_{i-1})}{(\mathrm{d}\mathbf{P}_0(Z_i|Z_{i-1}))^2}\Big|\{v^{(j)}\}_{1\leq j\leq i}\right]. \tag{6.32}$$

Moreover, let $\mathcal{V}_\theta^k = \left\{v^{(1)}, \ldots, v^{(k)} \in (\mathcal{S}^{d-1})^k \;:\; \forall i \in [k],\; d\cdot\langle v^{(i)}, \theta\rangle^2 \leq \tau_i\right\}$. Then,

$$\mathbb{E}_{\mathbf{P}_0}\left[\frac{\mathrm{d}\mathbf{P}_{u_1}(Z_k; A_{u_1}^k)\mathrm{d}\mathbf{P}_{u_2}(Z_k; A_{u_2}^k)}{\left(\mathrm{d}\mathbf{P}_0(Z_k)\right)^2}\right] \leq \sup_{v^{(1)},\ldots,v^{(k)}\in\mathcal{V}_{u_1}^k\cap\mathcal{V}_{u_2}^k}\;\prod_{i=1}^k g_i\left(u_1, u_2; \{v^{(j)}\}_{1\leq j\leq i}\right). \tag{6.33}$$

With a bit of computation, one can make the above bound more explicit (see Proposition C.4):

$$\mathbb{E}_{\mathbf{P}_0}\left[\frac{\mathrm{d}\mathbf{P}_{u_1}(Z_T; A_{u_1}^T)\mathrm{d}\mathbf{P}_{u_2}(Z_T; A_{u_2}^T)}{\left(\mathrm{d}\mathbf{P}_0(Z_T)\right)^2}\right] \leq e^{\lambda^2|\langle u_1,u_2\rangle|\sum_{i=1}^T\tau_i + \frac{\lambda^2}{d}(\sum_{i=1}^T\tau_i)^2}. \tag{6.34}$$

To wrap up, we will again take $\mathcal{P}$ to be the uniform measure on the sphere, which satisfies the following moment bound.

Lemma 6.3. *Let $\mathcal{P}$ be the uniform measure on the sphere $\mathcal{S}^{d-1}$. Then for all $\lambda \geq 0$ and $v \in \mathcal{S}^{d-1}$,*

$$\mathbb{E}_{\theta\sim\mathcal{P}}\left[e^{\lambda|\langle\theta,v\rangle|}\right] \leq e^{4\lambda^2/d + \lambda\sqrt{2/d}}. \tag{6.35}$$

Proof of Lemma 6.3. Write $\mathbb{E}[e^{\sqrt{d}\lambda|\langle\theta,v\rangle|}] = e^{\lambda\sqrt{2}}\cdot\mathbb{E}[e^{\lambda(\sqrt{d}|\langle\theta,v\rangle|-\sqrt{2}\lambda)}]$. Since $\sqrt{d}(|\langle\theta, v\rangle| - \sqrt{2})$ satisfies the sub-Gaussian tail bound from Lemma 4.4, a standard conversion from tail bounds to moment generating functions (see, e.g., Lemma 1.5 in [69]) yields that $\mathbb{E}[e^{\lambda\cdot\sqrt{d}(|\langle\theta,v\rangle|-\sqrt{2})}] \leq e^{4\lambda^2}$. Replacing λ with $\lambda/\sqrt{d}$ concludes the proof. □

Now, we take an expectation of the bound in Equation (6.34) over $u_1, u_2 \sim \mathcal{P}$, and apply the sub-Gaussian bound of Lemma 6.3. Further, appropriately choosing the parameters $\tau_1, \cdots, \tau_T$, and applying Proposition 6.1 leads to the following result.

Proposition 6.4. *There exists a absolute constant $c_1 > 1$ such that for all $\lambda > 2$ we have*

$$\|\mathbf{P}_0 - \overline{\mathbf{P}}\|_{\mathrm{TV}} \leq \frac{\sqrt{2}\,(c_1\lambda)^T}{d^{1/4}} \cdot \left(\sqrt{\log \frac{d}{(c_1\lambda)^T}} + 4 \right). \tag{6.36}$$

The proof of the above Proposition is detailed in Appendix D.4. We now state and prove our main technical result for detecting a large eigenvalue.

Theorem 6.5 (Technical Statement of Theorem 2.3)**.** *For any $\delta_0 \in (0,1)$ and $\lambda \geq K_d + 4d^{-1/2}\sqrt{\log(1/\delta_0)}$, there exist two distributions $\mathcal{D}_0$ and $\mathcal{D}_1$ on $\mathbb{S}^{d\times d}$ such that*

1. *$\|M\| \leq K_d + 2d^{-1/2}\sqrt{\log(1/\delta_0)}$ $\mathcal{D}_0$-almost surely,*
2. *$\|M\| \geq \lambda - 2d^{-1/2}\sqrt{\log(1/\delta_0)}$ $\mathcal{D}_1$-almost surely,*

and, for any T-query algorithm which outputs a binary test $\psi \in \{0,1\}$,

$$\mathbb{P}_{M\sim\mathcal{D}_0,\mathsf{Alg}}[\psi = 1] + \mathbb{P}_{M\sim\mathcal{D}_1,\mathsf{Alg}}[\psi = 0] \geq 1 - \frac{\sqrt{2}\,(c_1\lambda)^T}{d^{1/4}} \cdot \left(\sqrt{\log \frac{d}{(c_1\lambda)^T}} + 4 \right) - 3\delta_0. \tag{6.37}$$

Moreover, $\mathcal{D}_1$-almost surely, $M \in M_\gamma$ for $\gamma(\lambda,\delta)$ as defined in Proposition 3.2. Here, c_1 is the absolute constant of Proposition 6.4.

Proof of Theorem 6.5. Let $\overline{A}$ denote the event of Proposition 3.2 marginalized over $\theta \sim \mathcal{P}$. This event holds with probability at least $1 - 2\delta_0$ under $\overline{\mathbb{P}}$. The proof of Theorem 6.5 also show that, with probability $1 - \delta_0$ under $\mathbb{P}_0$, $\|W\| \leq K_d + 2\sqrt{\log(1/\delta_0)}$; we denote this event A_0. Then, we have that

$$\begin{aligned}
\mathbb{P}_0[\psi = 1|A_0] + \overline{\mathbb{P}}\left[\psi = 0|\overline{A}\right] &\geq \mathbb{P}_0[\psi = 1] + \overline{\mathbb{P}}[\psi = 0] - 3\delta_0 \\
&\overset{(i)}{\geq} 1 - \|\mathbf{P}_0 - \overline{\mathbf{P}}\|_{\mathrm{TV}} - 3\delta_0 \\
&\overset{(ii)}{\geq} 1 - \frac{\sqrt{2}\left(c_2\lambda^2\right)^{T/2}}{d^{1/4}} \cdot \left(\sqrt{\log \frac{d}{(\gamma'' c_2)^T}} + 4 \right) - 3\delta_0.
\end{aligned}$$

where (i) follows from a standard hypothesis-testing inequality (Theorem 2.2 in [81]) and the fact that ψ is a measurable function of Z_T, and (ii) is Lemma 6.4 above. Finally, on $\overline{A}$, $M \in \mathcal{M}_\gamma$ for $\gamma = \gamma(\lambda,\delta)$ defined in Propostion 3.2. Finally take $\mathcal{D}_0 := \mathbb{P}_0[\cdot|A_0]$ and $\mathcal{D}_1 := \overline{\mathbb{P}}[\cdot|\overline{A}]$. □

7 Conclusion

This paper established a fundamental separation between the first-order query complexity for optimizing strict-saddle and truly convex objectives. We demonstrated the separation by establishing a query complexity lower bound of $\Omega(\log d)$ in the easy regime of rank-one PCA where the leading and second eigenvalues are well separated. An exciting direction for future work is to attempt to prove

lower bounds for the "hard" regime of rank-one PCA, where the gap between the leading and second eigenvalues, $1-\gamma$, is arbitrarily small. Ideally one would show $\Omega\left(\frac{\log(d/\epsilon)}{\sqrt{1-\gamma}}\right)$ gap-dependent and $\Omega\left(\frac{\log d}{\sqrt{\epsilon}}\right)$ query complexity lower bounds, thereby matching the randomized block-Krylov methods from Musco and Musco [56].

We believe that establishing such a lower bound would entail numerous technical challenges. As remarked at the end of Theorem 5.4, establishing this bound for the deformed Wigner model would requires obtaining sharp control on the rate of information accumulation, at the so-called "identifiability threshold", where λ tends to 1[4]. One would also need to invoke non-trivial machinery from random matrix theory [33] to sharpen Proposition 3.2 in order to establish that $v_1(M)$ correlates with θ when $\lambda \in [1, K_d]$.

Another direction for future research would be to understand if the classical lower bounds against Krylov methods for *convex* optimization [60] hold in the stronger gradient-query model as well. This would resolve objections raised in the literature that these lower bounds place unduly strong assumptions on the class of gradient-methods considered, and in particular, do not apply when the starting point of the algorithm may be randomized. More broadly, understanding the query complexity of both non-convex and convex optimization would serve to elucidate the connections between sequential optimization and adaptive estimation.

Acknowledgements

We thank Chi Jin and Darren (Tianyi) Lin for their immensely useful feedback. Max Simchowitz is supported by an NSF GRFP fellowship. Benjamin Recht is generously supported by NSF award CCF-1359814, ONR awards N00014-14-1-0024 and N00014-17-1-2191, the DARPA Fundamental Limits of Learning Program, a Sloan Research Fellowship, and a Google Faculty Award.

References

[1] Dimitris Achlioptas and Cristopher Moore. Random k-sat: Two moments suffice to cross a sharp threshold. *SIAM Journal on Computing*, 36(3):740–762, 2006.

[2] Louigi Addario-Berry, Nicolas Broutin, Luc Devroye, Gábor Lugosi, et al. On combinatorial testing problems. *The Annals of Statistics*, 38(5):3063–3092, 2010.

[3] Alekh Agarwal, Martin J Wainwright, Peter L Bartlett, and Pradeep K Ravikumar. Information-theoretic lower bounds on the oracle complexity of convex optimization. In *Advances in Neural Information Processing Systems*, pages 1–9, 2009.

[4] Naman Agarwal, Zeyuan Allen-Zhu, Brian Bullins, Elad Hazan, and Tengyu Ma. Finding approximate local minima for nonconvex optimization in linear time. *arXiv preprint arXiv:1611.01146*, 2016.

[5] Greg W Anderson, Alice Guionnet, and Ofer Zeitouni. *An introduction to random matrices*, volume 118. Cambridge university press, 2010.

[4] It is a non-obvious fact from random matrix theory that the deformation is detectable even when $\lambda < \mathbb{E}[\|W\|]$. Indeed $\|W + \lambda\theta\theta^\top\|$ concentrates around $\lambda + 1/\lambda$ for $\lambda \geq 1$. This implies that $\|W + \lambda\theta\theta^\top\| > \|W\|$, even when $\lambda < 2$.

[6] Anurag Anshu, Naresh B Goud, Rahul Jain, Srijita Kundu, and Priyanka Mukhopadhyay. Lifting randomized query complexity to randomized communication complexity. *arXiv preprint arXiv:1703.07521*, 2017.

[7] Ery Arias-Castro, Emmanuel J Candes, and Mark A Davenport. On the fundamental limits of adaptive sensing. *IEEE Transactions on Information Theory*, 59(1):472–481, 2013.

[8] Yossi Arjevani, Shai Shalev-Shwartz, and Ohad Shamir. On lower and upper bounds in smooth and strongly convex optimization. *Journal of Machine Learning Research*, 17(126):1–51, 2016.

[9] Yossi Arjevani and Ohad Shamir. Oracle complexity of second-order methods for finite-sum problems. *arXiv preprint arXiv:1611.04982*, 2016.

[10] Sanjeev Arora, Rong Ge, Ravindran Kannan, and Ankur Moitra. Computing a nonnegative matrix factorization–provably. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 145–162. ACM, 2012.

[11] Sanjeev Arora, Rong Ge, Tengyu Ma, and Ankur Moitra. Simple, efficient, and neural algorithms for sparse coding. In *COLT*, pages 113–149, 2015.

[12] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*, 2014.

[13] Maria-Florina Balcan, Yingyu Liang, Le Song, David Woodruff, and Bo Xie. Communication efficient distributed kernel principal component analysis. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 725–734. ACM, 2016.

[14] Afonso S Bandeira, Ramon van Handel, et al. Sharp nonasymptotic bounds on the norm of random matrices with independent entries. *The Annals of Probability*, 44(4):2479–2506, 2016.

[15] Srinadh Bhojanapalli, Anastasios Kyrillidis, and Sujay Sanghavi. Dropping convexity for faster semi-definite optimization. In *29th Annual Conference on Learning Theory*, pages 530–582, 2016.

[16] Srinadh Bhojanapalli, Behnam Neyshabur, and Nati Srebro. Global optimality of local search for low rank matrix recovery. In *Advances in Neural Information Processing Systems*, pages 3873–3881, 2016.

[17] Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. *Concentration inequalities: A nonasymptotic theory of independence*. Oxford university press, 2013.

[18] Nicolas Boumal. Nonconvex phase synchronization. *SIAM Journal on Optimization*, 26(4):2355–2377, 2016.

[19] Christos Boutsidis, David P Woodruff, and Peilin Zhong. Optimal principal component analysis in distributed and streaming models. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing*, pages 236–249. ACM, 2016.

[20] Stephen Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge university press, 2004.

[21] Sébastien Bubeck et al. Convex optimization: Algorithms and complexity. *Foundations and Trends® in Machine Learning*, 8(3-4):231–357, 2015.

[22] Emmanuel J Candes, Xiaodong Li, and Mahdi Soltanolkotabi. Phase retrieval via wirtinger flow: Theory and algorithms. *IEEE Transactions on Information Theory*, 61(4):1985–2007, 2015.

[23] Yair Carmon and John C Duchi. Gradient descent efficiently finds the cubic-regularized non-convex newton step. *arXiv preprint arXiv:1612.00547*, 2016.

[24] Rui M Castro et al. Adaptive sensing performance lower bounds for sparse signal detection and support estimation. *Bernoulli*, 20(4):2217–2246, 2014.

[25] Rui M Castro and Ervin Tánczos. Adaptive compressed sensing for support recovery of structured sparse sets. *IEEE Transactions on Information Theory*, 63(3):1535–1554, 2017.

[26] Xi Chen, Adityanand Guntuboyina, and Yuchen Zhang. On bayes risk lower bounds. *Journal of Machine Learning Research*, 17:1–58, 2016.

[27] Kenneth L Clarkson and David P Woodruff. Low rank approximation and regression in input sparsity time. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 81–90. ACM, 2013.

[28] Kenneth L Clarkson and David P Woodruff. Low-rank psd approximation in input-sparsity time. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2061–2072. SIAM, 2017.

[29] Thomas M Cover and Joy A Thomas. *Elements of information theory*. John Wiley & Sons, 2012.

[30] Imre Csiszár. A class of measures of informativity of observation channels. *Periodica Mathematica Hungarica*, 2(1-4):191–213, 1972.

[31] James W Demmel. *Applied numerical linear algebra*. SIAM, 1997.

[32] John C Duchi and Martin J Wainwright. Distance-based and continuum fano inequalities with applications to statistical estimation. *arXiv preprint arXiv:1311.2669*, 2013.

[33] Delphine Féral and Sandrine Péché. The largest eigenvalue of rank one deformation of large wigner matrices. *Communications in mathematical physics*, 272(1):185–228, 2007.

[34] James R Fienup. Phase retrieval algorithms: a comparison. *Applied optics*, 21(15):2758–2769, 1982.

[35] Dan Garber, Elad Hazan, Chi Jin, Sham M Kakade, Cameron Musco, Praneeth Netrapalli, and Aaron Sidford. Faster eigenvector computation via shift-and-invert preconditioning. In *Proceedings of the 33nd International Conference on Machine Learning, ICML*, pages 2626–2634, 2016.

[36] Dan Garber, Elad Hazan, and Tengyu Ma. Online learning of eigenvectors. In *ICML*, pages 560–568, 2015.

[37] Dan Garber, Ohad Shamir, and Nathan Srebro. Communication-efficient algorithms for distributed stochastic principal component analysis. *arXiv preprint arXiv:1702.08169*, 2017.

[38] Aurélien Garivier and Emilie Kaufmann. Optimal best arm identification with fixed confidence. In *Proceedings of the 29th Conference On Learning Theory (to appear)*, 2016.

[39] Rong Ge, Furong Huang, Chi Jin, and Yang Yuan. Escaping from saddle points-online stochastic gradient for tensor decomposition. In *COLT*, pages 797–842, 2015.

[40] Rong Ge, Jason D Lee, and Tengyu Ma. Matrix completion has no spurious local minimum. In *Advances in Neural Information Processing Systems*, pages 2973–2981, 2016.

[41] Adityanand Guntuboyina. Lower bounds for the minimax risk using f-divergences, and applications. *IEEE Transactions on Information Theory*, 57(4):2386–2399, 2011.

[42] Roger A Horn and Charles R Johnson. *Matrix analysis*. Cambridge university press, 2012.

[43] Kevin G Jamieson, Robert Nowak, and Ben Recht. Query complexity of derivative-free optimization. In *Advances in Neural Information Processing Systems*, pages 2672–2680, 2012.

[44] Majid Janzamin, Hanie Sedghi, and Anima Anandkumar. Beating the perils of non-convexity: Guaranteed training of neural networks using tensor methods. *arXiv preprint arXiv:1506.08473*, 2015.

[45] Chi Jin, Rong Ge, Praneeth Netrapalli, Sham M Kakade, and Michael I Jordan. How to escape saddle points efficiently. *arXiv preprint arXiv:1703.00887*, 2017.

[46] Ian Jolliffe. *Principal component analysis*. Wiley Online Library, 2002.

[47] Olav Kallenberg. *Foundations of modern probability*. Springer Science & Business Media, 2006.

[48] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.

[49] Daniel D Lee and H Sebastian Seung. Learning the parts of objects by non-negative matrix factorization. *Nature*, 401(6755):788–791, 1999.

[50] Jason D Lee, Max Simchowitz, Michael I Jordan, and Benjamin Recht. Gradient descent only converges to minimizers. In *Conference on Learning Theory*, pages 1246–1257, 2016.

[51] Marc Lelarge and Léo Miolane. Fundamental limits of symmetric low-rank matrix estimation. *arXiv preprint arXiv:1611.03888*, 2016.

[52] Yi Li, Huy L Nguyên, and David P Woodruff. On sketching matrix norms and the top singular vector. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1562–1581. Society for Industrial and Applied Mathematics, 2014.

[53] Yi Li and David P Woodruff. On approximating functions of the singular values in a stream. *arXiv preprint arXiv:1604.08679*, 2016.

[54] Friedrich Liese. Phi-divergences, sufficiency, bayes sufficiency, and deficiency. *Kybernetika*, 48(4):690–713, 2012.

[55] Katta G Murty and Santosh N Kabadi. Some np-complete problems in quadratic and nonlinear programming. *Mathematical programming*, 39(2):117–129, 1987.

[56] Cameron Musco and Christopher Musco. Randomized block krylov methods for stronger and faster approximate singular value decomposition. In *Advances in Neural Information Processing Systems*, pages 1396–1404, 2015.

[57] Jelani Nelson. Time lower bounds for nonadaptive turnstile streaming algorithms. In *Banff International Research Station for Mathematical Innovation and Discovery (BIRS) Workshop Lecture Videos*. Banff International Research Station for Mathematical Innovation and Discovery, 2014.

[58] Jelani Nelson, Huy L Nguyn, and David P Woodruff. On deterministic sketching and streaming for sparse recovery and norm estimation. *Linear Algebra and its Applications*, 441:152–167, 2014.

[59] Jelani Nelson, Jakub Pachocki, and Zhengyu Wang. Optimal lower bounds for universal relation, samplers, and finding duplicates. *arXiv preprint arXiv:1703.08139*, 2017.

[60] Arkadii Nemirovskii, David Borisovich Yudin, and Edgar Ronald Dawson. Problem complexity and method efficiency in optimization. 1983.

[61] Yu Nesterov. Introductory lectures on convex programming volume i: Basic course. 1998.

[62] Praneeth Netrapalli, Prateek Jain, and Sujay Sanghavi. Phase retrieval using alternating minimization. In *Advances in Neural Information Processing Systems*, pages 2796–2804, 2013.

[63] Andrew Y Ng, Michael I Jordan, et al. On spectral clustering: Analysis and an algorithm.

[64] Bruno A Olshausen and David J Field. Sparse coding with an overcomplete basis set: A strategy employed by v1? *Vision research*, 37(23):3311–3325, 1997.

[65] Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. The pagerank citation ranking: Bringing order to the web. Technical report, Stanford InfoLab, 1999.

[66] Eric Price and David P Woodruff. Lower bounds for adaptive sparse recovery. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 652–663. Society for Industrial and Applied Mathematics, 2013.

[67] Ilya P Razenshteyn, Zhao Song, and David P Woodruff. Weighted low rank approximations with provable guarantees. In *STOC*, pages 250–263, 2016.

[68] Jasson DM Rennie and Nathan Srebro. Fast maximum margin matrix factorization for collaborative prediction. In *Proceedings of the 22nd international conference on Machine learning*, pages 713–719. ACM, 2005.

[69] Philippe Rigollet. 18. s997: High dimensional statistics. 2015.

[70] Ohad Shamir. Fundamental limits of online and distributed algorithms for statistical learning and estimation. In *Advances in Neural Information Processing Systems*, pages 163–171, 2014.

[71] Ohad Shamir. Fast stochastic algorithms for svd and pca: Convergence properties and convexity. *arXiv preprint arXiv:1507.08788*, 2015.

[72] Max Simchowitz, Kevin Jamieson, and Benjamin Recht. The simulator: Understanding adaptive sampling in the moderate-confidence regime. *arXiv preprint arXiv:1702.05186*, 2017.

[73] Marta Soare, Alessandro Lazaric, and Rémi Munos. Best-arm identification in linear bandits. In *Advances in Neural Information Processing Systems*, pages 828–836, 2014.

[74] Zhao Song, David P Woodruff, and Peilin Zhong. Low rank approximation with entrywise l1 -norm error. *arXiv preprint arXiv:1611.00898*, 2016.

[75] Daniel A Spielman. Spectral graph theory and its applications. In *Foundations of Computer Science, 2007. FOCS'07. 48th Annual IEEE Symposium on*, pages 29–38. IEEE, 2007.

[76] Jacob Steinhardt and John C Duchi. Minimax rates for memory-bounded sparse linear regression. In *COLT*, pages 1564–1587, 2015.

[77] Jacob Steinhardt, Gregory Valiant, and Stefan Wager. Memory, communication, and statistical queries. In *Electronic Colloquium on Computational Complexity (ECCC)*, volume 22, pages 1–2, 2015.

[78] Ju Sun, Qing Qu, and John Wright. When are nonconvex problems not scary? *arXiv preprint arXiv:1510.06096*, 2015.

[79] Ju Sun, Qing Qu, and John Wright. A geometric analysis of phase retrieval. In *Information Theory (ISIT), 2016 IEEE International Symposium on*, pages 2379–2383. IEEE, 2016.

[80] Ju Sun, Qing Qu, and John Wright. Complete dictionary recovery over the sphere i: Overview and the geometric picture. *IEEE Transactions on Information Theory*, 63(2):853–884, 2017.

[81] Alexandre B Tsybakov. Introduction to nonparametric estimation. revised and extended from the 2004 french original. translated by vladimir zaiats, 2009.

[82] Stephen Tu, Ross Boczar, Max Simchowitz, Mahdi Soltanolkotabi, and Benjamin Recht. Low-rank solutions of linear matrix equations via procrustes flow. *arXiv preprint arXiv:1507.03566*, 2015.

[83] Roman Vershynin. Introduction to the non-asymptotic analysis of random matrices. *arXiv preprint arXiv:1011.3027*, 2010.

[84] Andrew Chi-Chin Yao. Probabilistic computations: Toward a unified measure of complexity. In *Foundations of Computer Science, 1977., 18th Annual Symposium on*, pages 222–227. IEEE, 1977.

[85] Qinqing Zheng and John Lafferty. A convergent gradient descent algorithm for rank minimization and semidefinite programming from random linear measurements. In *Advances in Neural Information Processing Systems*, pages 109–117, 2015.

[86] Yunhong Zhou, Dennis Wilkinson, Robert Schreiber, and Rong Pan. Large-scale parallel collaborative filtering for the netflix prize. In *International Conference on Algorithmic Applications in Management*, pages 337–348. Springer, 2008.

A Proof of Proposition 3.2

In this section, we prove the validity of our reduction. To prove the strongest lower bounds, we want to allow λ to be arbitrarily close to $\lambda \approx K_d = 2+o(1)$. This precludes using coarser arguments, e.g., bounding the $\|W\|_{op}$ and applying the Davis-Kahan Sine theorem. We start with a deterministic result which sharpens Davis-Kahan in our particular case of interest:

Lemma A.1. *Let $M = \lambda\theta\theta^\top + W$, where $\theta \in \mathcal{S}^{d-1}$ and W is an arbitrary symmetric matrix. Set $\gamma := \frac{\|W\|}{\theta^\top M\theta}$. Then, for any $\epsilon \le \gamma$, if $w^\top Mw \ge (1-\epsilon)\theta^\top M\theta$, then*

$$|\langle w, \theta\rangle| \ge F(\epsilon,\gamma) := \sqrt{\left(\frac{\gamma}{1-\gamma}\right)^2 + 1 - \frac{\epsilon}{1-\gamma}} - \frac{\gamma}{1-\gamma}.$$

Moroever, $\lambda_1(M) \ge \theta^\top M\theta$, and $\lambda_i(M) \le \|W\|$ for all $i \ge 2$.

We now apply above lemma with $W/\sqrt{d}$. To conclude the proof of Proposition 3.2, it suffices to verify that the following two conditions hold with probability $1-\delta$,

$$\|W\| \le K_d + 2\sqrt{\log(1/\delta)/d} \quad \text{and} \quad \theta^\top M\theta - \lambda = \theta^\top W\theta \ge -2\sqrt{\log(1/\delta)/d}. \tag{1.38}$$

To this end, we use an alternate characterization of the GOE:

Lemma A.2. *$W \sim \mathrm{GOE}(d)$ has the distribution $\frac{1}{\sqrt{2}}(X + X^\top)$, where $X \in \mathbb{R}^{d\times d}$ is a matrix with i.i.d. $\mathcal{N}(0,1)$ entries.*

Proof. Let $\widetilde{W} := \frac{1}{\sqrt{2}}(X + X^\top)$, where X has i.i.d. standard normal entries. Then $\widetilde{W}$ is symmetric, the entries $\widetilde{W}_{i,j}$ for $i \le j$ are independent centered Gaussians with $\mathbb{E}[\widetilde{W}_{ii}^2 = \mathbb{E}[(\sqrt{2}X_{ii})^2] = 2$, while for $i \ne j$, $\mathbb{E}[\widetilde{W}_{ii}^2 = \frac{1}{2}\mathbb{E}[(X_{ij} + X_{ji})^2] = 1$ since X_{ij} and X_{ji} are independent. □

Note then that if $X \in \mathbb{R}^{d\times d}$ is a matrix with i.i.d. $\mathcal{N}(0,1)$, then $f_1(X) := \theta^\top \frac{1}{\sqrt{2}}(X + X^\top)\theta$ has the same distribution as $\theta^\top W\theta$, and $f_2(X) := \|\frac{1}{\sqrt{2}}(X + X^\top)\|$ has the same distribution as $\|W\|$. The following lemma shows that f_1 and f_2 are both $\sqrt{2}$-Lipschitz.

Lemma A.3. *Let v be a unit vector. Then the mappings $X \mapsto v^\top \frac{1}{\sqrt{2}}(X + X^\top)v$, $X \mapsto \|\frac{1}{\sqrt{2}}(X + X^\top)\|$ are $\sqrt{2}$-Lipschitz.*

Proof. By Cauchy Schwartz, we have for any $v, w \in \mathcal{S}^{d-1}$

$$|v^\top \frac{1}{\sqrt{2}}(X + X^\top)w| = \sqrt{2}|v^\top Xw| \le \sqrt{2}\cdot\|X\|_F\|wv^\top\|_F = \sqrt{2}\|X\|_F.$$

Since the Lipschitz constant of a linear map is equal to its operator norm, we have that $x \mapsto v^\top \frac{1}{\sqrt{2}}(X + X^\top)w$ is $\sqrt{2}$-Lipschitz. The first part of the lemma follows by taking $v = w$. The second point follows by noting that $\|\frac{1}{\sqrt{2}}(X + X^\top)\|_2 = \sup_{v,w\in\mathcal{S}^{d-1}} v^\top \frac{1}{\sqrt{2}}(X + X^\top)w$, and the supremum of a collection of L-Lipschitz functions is L-Lipschitz. □

Equation (1.38) now follows by applying the following concentration inequality to the functions $f_1(X)$ and $f_2(X)$:

Lemma A.4 (Tsirelson-Ibgragimov-Sudakov, Theorem 5.5 in [17])**.** *Let f be a L-Lipschitz function and let X be a standard Gaussian vector. Then,*

$$\mathbb{P}[f(X) \ge \mathbb{E}[f(X)] + t] \vee \mathbb{P}[f(X) - \mathbb{E}[f(X)] \le -t] \le e^{-t^2/2L^2}.$$

A.1 Proof of Lemma A.1

Recall that $\theta \in \mathcal{S}^{d-1}$ Any vector $w \in \mathcal{S}^{d-1}$ can be written as $\alpha\theta + \sqrt{1-\alpha^2}v$, where $v \perp \theta$ and $v \in \mathcal{S}^{d-1}$. Note then that $\alpha = \langle w, \theta\rangle$. By replacing v with $-v$, we may assume without loss of generality that $\alpha \geq 0$. For ease of notation, set $K_2 = \theta^\top M\theta$ and $K_1 = \|W\|$. Then, $w^\top Mw$ can be written as

$$\begin{aligned}(\alpha\theta + \sqrt{1-\alpha^2}v)^\top M(\alpha\theta + \sqrt{1-\alpha^2}v) &= \alpha^2\theta^\top M\theta + (1-\alpha^2)v^\top Mv + 2\alpha\sqrt{1-\alpha^2}v^\top M\theta \\ &\leq \alpha^2\theta^\top M\theta + (1-\alpha^2)K_1 + 2|\alpha|K_1.\end{aligned}$$

Thus, if $w^\top Mw \geq (1-\epsilon)\|M\| \geq (1-\epsilon)K_2$, then

$$(1-\epsilon-\alpha^2)K_2 - (1-\alpha^2)K_1 - 2\alpha K_1 \leq 0.$$

Letting $\gamma = K_1/K_2$, we have

$$\alpha^2(1-\gamma) - (1-\epsilon-\gamma) + 2\alpha\gamma \geq 0.$$

Thus, solving the quadratic inequality and taking the positive part,

$$\begin{aligned}\alpha &\geq \frac{-2\gamma + \sqrt{4\gamma^2 + 4(1-\gamma)(1-\epsilon-\gamma)}}{2(1-\gamma)} \\ &= \sqrt{(\frac{\gamma}{1-\gamma})^2 + 1 - \frac{\epsilon}{1-\gamma}} - \frac{\gamma}{1-\gamma}.\end{aligned}$$

The second statement in Lemma A.1 follows since $\lambda_1(M) \geq \theta^\top M\theta = \lambda + \theta^\top W\theta$ and $\lambda \geq 0$, and the third statement follows from eigenvalue interlacing (e.g., Corollary 4.3.9 in [42]).

A.2 Behavior of $F(\epsilon, \gamma)$

The following lemma describes the behavior of $F(\epsilon, \gamma)$ is both the small- and large-γ regimes:

Lemma A.5. $\lim_{\gamma\to 0} F(\epsilon, \gamma) = \sqrt{1-\epsilon}$, *and for any* $\epsilon \in (0, 1-\gamma)$,

$$F(\epsilon, \gamma) \geq \frac{1}{2\sqrt{2}} \min\left\{\sqrt{1 - \frac{\epsilon}{1-\gamma}}, \frac{1-\gamma-\epsilon}{\gamma}\right\}. \tag{1.39}$$

Proof. The computation of $\lim_{\gamma\to 0} F(\epsilon, \gamma)$ is clear from the definition of $F(\epsilon, \gamma)$. For the second statement, we see that Taylors theorem implies $\sqrt{a^2+x} - a \geq \frac{x}{2\sqrt{a^2+x}} \geq \frac{x}{2\sqrt{2}} \min\{1/a, 1/\sqrt{x}\}$. This entails

$$\begin{aligned}F(\epsilon, \gamma) &\geq \frac{1}{2}(1 - \frac{\epsilon}{1-\gamma}) \cdot \left(\sqrt{\left(\frac{\gamma}{1-\gamma}\right)^2 + 1 - \frac{\epsilon}{1-\gamma}}\right)^{-1} \\ &\geq \frac{1}{2\sqrt{2}}\left(1 - \frac{\epsilon}{1-\gamma}\right) \cdot \left(\max\left\{\left(\frac{\gamma}{1-\gamma}\right)^2, 1 - \frac{\epsilon}{1-\gamma}\right\}\right)^{-1/2} \\ &= \frac{1}{2\sqrt{2}}\left(1 - \frac{\epsilon}{1-\kappa}\right) \cdot \min\left\{\frac{1-\gamma}{\gamma}, \left(1 - \frac{\epsilon}{1-\gamma}\right)^{-1/2}\right\} \\ &= \frac{1}{2\sqrt{2}} \min\left\{\sqrt{1 - \frac{\epsilon}{1-\kappa}}, \frac{1-\gamma-\epsilon}{\gamma}\right\}.\end{aligned}$$

□

B Proof Lemma 3.4

Recall the definition $\Sigma_i := P_{i-1}(I_d + v^{(i)}v^{(i)\top})P_{i-1}$, and that $\mathcal{F}_{i-1}$ is the σ-algebra generated by $v^{(1)}, \widetilde{w}^{(1)}, \dots, \widetilde{v}^{(i-1)}, \widetilde{w}^{(i-1)}$. Since our algorithm is deterministic, $v^{(i)}$ is $\mathcal{F}_{i-1}$ measurable. It then suffices to show that

$$\widetilde{w}^{(j)} = P_{j-1}Wv^{(j)}\big|\mathcal{F}_i \sim \mathcal{N}(0, \Sigma_i) \ . \tag{2.40}$$

Recall from Section 3.1 that Σ_i is degenerate, so we understand $\mathcal{N}(0, \frac{1}{d}\Sigma_i)$ as a normal distribution absolutely continuous with respect to the Lebesque measure supported on $(\ker P_{i-1})^\perp$. Note that $\widetilde{w}^{(i)}$ is conditionally independent of $\widetilde{w}^{(1)}, \dots, \widetilde{w}^{(i-1)}$ given $v^{(1)}, \dots, v^{(i-1)}, v^{(i)}$. Consequently, the conditional distribution of $\widetilde{w}^{(i)}$ given $\mathcal{F}_i$ can be computed as if the queries $v^{(1)}, \dots, v^{(i)}$ were fixed in advanced.

Hence, throughout, we shall assume that $v^{(1)}, \dots, v^{(i)}$ are deterministic, and consider the joint distribution of $\widetilde{w}^{(1)}, \dots, \widetilde{w}^{(i-1)}, \widetilde{w}^{(i)}$. We will show that $\widetilde{w}^{(i)}$ is independent of $\widetilde{w}^{(1)}, \dots, \widetilde{w}^{(i-1)}$, and that its marginal is $\mathcal{N}(0, \frac{1}{d}\Sigma_i)$. Since the map $W \mapsto P_{j-1}Wv^{(j)}$ is linear maps, $\widetilde{w}^{(1)}, \dots, \widetilde{w}^{(i)}$ are jointly Gaussian with mean zero. Thus, it suffices to show that 1) the (marginal) covariance of $\widetilde{w}^{(i)}$ is Σ_i and, 2) the covariance between $\widetilde{w}^{(i)}$ and $\widetilde{w}^{(j)}$ for $j \neq i$ is 0. The covariances are computed as

$$\mathbb{E}\left[\widetilde{w}^{(j)}\widetilde{w}^{(j)\top}\right] = \mathbb{E}\left[(P_{i-1}Wv^{(i)})(P_{j-1}Wv^{(j)})^\top\right] = P_{i-1}\mathbb{E}\left[Wv^{(i)}v^{(j)\top}W\right]P_{j-1} \ , \tag{2.41}$$

and we compute the inner term with the following lemma.

Lemma B.1. *For any $v^{(i)}, v^{(j)}$, one has*

$$\mathbb{E}\left[Wv^{(i)}v^{(j)\top}W\right] = v^{(j)}v^{(i)\top} + \langle v^{(i)}, v^{(j)}\rangle I \ . \tag{2.42}$$

For $v^{(i)} = v^{(j)}$, Equations (2.41) and (2.42) immediately imply $\widetilde{w}^{(i)}$ has covariance Σ_i. Moreover, for $j < i$, we have

$$P_{i-1}\mathbb{E}\left[Wv^{(i)}v^{(j)\top}W\right]P_{j-1} \overset{(i)}{=} P_{i-1}v^{(i)}v^{(j)\top}P_{j-1} \overset{(ii)}{=} 0 \tag{2.43}$$

where (i) holds from Lemma B.1 and the fact that $\langle v^{(i)}, v^{(j)}\rangle = 0$ (since $v^{(1)}, \dots, v^{(i)}$ are assumed to be orthogonal), and (ii) holds since $v^{(j)} \in \ker(P_{i-1})$, as P_{i-1} projects onto the complement of $\{v^{(1)}, \dots, v^{(i-1)}\}$.

Proof of Lemma B.1. For $a \in \{1, \cdots, d\}$,

$$\begin{aligned}\mathbb{E}[Wv^{(i)}v^{(j)\top}W]_{aa} &= \mathbb{E}[W_{aa}^2]v_a^{(i)}v_a^{(j)} + \sum_{p\neq a}\mathbb{E}[W_{ap}^2]v_a^{(i)}v_a^{(j)} \\ = \ 2v_a^{(i)}v_a^{(j)} + \sum_{p\neq a} v_p^{(i)}v_p^{(j)} &= v_a^{(i)}v_a^{(j)} + \sum_{p} v_p^{(i)}v_p^{(j)} = v_a^{(i)}v_a^{(j)} + \langle v^{(i)}, v^{(j)}\rangle \ .\end{aligned}$$

Whereas for $a \neq b$,

$$\mathbb{E}[Wv^{(i)}v^{(j)\top}W]_{ab} \ = \ \mathbb{E}\left[\sum_{p,q} W_{ap}W_{bq}v_p^{(i)}v_q^{(j)}\right] = \sum_{p,q}\mathbb{E}[W_{ap}W_{bq}]v_p^{(i)}v_q^{(j)} \ . \tag{2.44}$$

W_{ap} and W_{bq} are independent unless $(a,p) = (b,q)$ or $(a,p) = (q,b)$. If $a \neq b$, then this means the only term in the above sum which is non zero is $p = b$ and $q = a$, which yields $v_a^{(j)}v_b^{(i)}$. □

C Information-Theoretic Computations

C.1 Supporting proofs for Section 4

We begin by stating a well-known computation (see, e.g., [29]) for KL divergence between two normal distribution:

Lemma C.1. *Let $X \sim \mathcal{N}(\mu_X, \Sigma)$ and $Y \sim \mathcal{N}(\mu_Y, \Sigma)$, where $\Sigma \succeq 0$ and $\mu_X - \mu_Y \in (\ker \Sigma)^\perp$, then*

$$\mathsf{KL}(\mathbf{P}_X, \mathbf{P}_Y) = \frac{1}{2}\|\mu_X - \mu_Y\|_\Sigma^2. \tag{3.45}$$

We now compute the mutual information between Z_k and θ, thereby proving Proposition 4.1:

Proof of Proposition 4.1. Recall that $\mathbf{P}_{u_0}$ and $\mathbf{P}_{u_1}$ denote the restriction of the measures $\mathbb{P}_{u_0}$ and $\mathbb{P}_{u_1}$ to events which are measurable with respect to Z_k. We then compute

$$\begin{aligned}
\mathsf{KL}(\mathbf{P}_{u_0}, \mathbf{P}_{u_1}) &= \mathbb{E}_{\mathbf{P}_{u_0}} \log \frac{\mathrm{d}\mathbf{P}_{u_0}(Z_k)}{\mathrm{d}\mathbf{P}_{u_1}(Z_k)} \\
&= \mathbb{E}_{\mathbf{P}_{u_0}} \sum_{i=1}^k \log \frac{\mathrm{d}\mathbf{P}_{u_0}((w^{(i)}, v^{(i)})|Z_{i-1})}{\mathrm{d}\mathbf{P}_{u_1}((w^{(i)}, v^{(i)})|Z_{i-1})} \\
&= \sum_{i=1}^k \mathbb{E}_{\mathbf{P}_{u_0}} \mathbb{E}_{\mathbf{P}_{u_0}} \left[\log \frac{\mathrm{d}\mathbf{P}_{u_0}(w^{(i)}|Z_{i-1})}{\mathrm{d}\mathbf{P}_{u_1}(w^{(i)}|Z_{i-1})} \Big| Z_{i-1} \right] \\
&\overset{(i)}{=} \sum_{i=1}^k \mathbb{E}_{\mathbf{P}_{u_0}} \left[\mathsf{KL}\left(\mathcal{N}\left(u_0^\top v^{(i)} \lambda P_{i-1} u_0, \frac{1}{d}\Sigma_i \right), \mathcal{N}\left(u_1^\top v^{(i)} \lambda P_{i-1} u_1, \frac{1}{d}\Sigma_i \right) \right) \right].
\end{aligned}$$

where (i) uses the fact that $v^{(i)}$ is Z_{i-1}-measurable, and uses the computation of the conditional distribution of $w^{(i)}$ from Lemma 3.4. We then compute

$$\begin{aligned}
&\mathsf{KL}\Big(\mathcal{N}\Big(u_0^\top v^{(i)} \lambda P_{i-1} u_0, \frac{1}{d}\Sigma_i \Big), \mathcal{N}\Big(u_1^\top v^{(i)} \lambda P_{i-1} u_1, \frac{1}{d}\Sigma_i \Big) \Big) \\
&\qquad \overset{(i)}{=} \frac{1}{2} \left\| u_0^\top v^{(i)} \lambda P_{i-1} u_0 - u_1^\top v^{(i)} \lambda P_{i-1} u_1 \right\|_{\frac{1}{d}\Sigma_i}^2 \\
&\qquad \overset{(ii)}{=} \frac{\lambda^2 d}{2} \left\| P_{i-1}(u_0^\top v^{(i)} u_0 - u_1^\top v^{(i)} u_1) \right\|_{\Sigma_i}^2 \\
&\qquad \overset{(iii)}{\leq} \frac{\lambda^2 d}{2} \left\| u_0^\top v^{(i)} u_0 - u_1^\top v^{(i)} u_1 \right\|^2 \\
&\qquad = \frac{\lambda^2 d}{2} \Big(\langle u_0, v^{(i)} \rangle^2 + \langle u_1, v^{(i)} \rangle^2 - 2\langle u_0, v^{(i)} \rangle \langle u_1, v^{(i)} \rangle \langle u_0, u_1 \rangle \Big),
\end{aligned} \tag{3.46}$$

where (i) uses Lemma C.1, and (ii) and (iii) use Equation (3.11) from Section (3.1). Hence, if $\mathcal{P}$

is isotropic, then $\mathbb{E}_{\theta\sim\mathcal{P}}[\theta\theta^\top] = \frac{1}{d}I$, so we can bound

$$\begin{aligned}
I(Z_k;\theta) &\overset{(i)}{\leq} \mathbb{E}_{u_0,u_1\sim\mathcal{P}}\mathsf{KL}(\mathbf{P}_{u_0},\mathbf{P}_{u_1}) \\
&\overset{(ii)}{\leq} \frac{\lambda^2 d}{2}\mathbb{E}_{u_0,u_1\sim\mathcal{P}}\sum_{i=1}^{k}\mathbb{E}_{\mathbf{P}_{u_0}}\left[\langle u_0,v^{(i)}\rangle^2 + \langle u_1,v^{(i)}\rangle^2 - 2v^{(i)\top}u_1u_1^\top u_0u_0^\top v^{(i)}\right] \\
&\overset{(iii)}{=} \frac{\lambda^2 d}{2}\mathbb{E}_{u_0\sim\mathcal{P}}\sum_{i=1}^{k}\mathbb{E}_{\mathbf{P}_{u_0}}\mathbb{E}_{u_1\sim\mathcal{P}}\left[\langle u_0,v^{(i)}\rangle^2 + \langle u_1,v^{(i)}\rangle^2 - 2v^{(i)\top}u_1u_1^\top u_0u_0^\top v^{(i)}\right] \\
&\overset{(iv)}{=} \frac{\lambda^2 d}{2}\mathbb{E}_{u_0\sim\mathcal{P}}\sum_{i=1}^{k}\left((1-\frac{2}{d})\mathbb{E}_{\mathbf{P}_{u_0}}\left[\langle u_0,v^{(i)}\rangle^2\right] + \frac{1}{d}\right) \\
&= \frac{\lambda^2}{2}\left\{k + (1-\frac{2}{d})\mathbb{E}_{u_0\sim\mathcal{P}}\mathbb{E}_{\mathbf{P}_{u_0}}\left[\sum_{i=1}^{k} d\langle u_0,v^{(i)}\rangle^2\right]\right\}.
\end{aligned}$$

where (i) uses convexity of the KL divergence [29], (ii) is Equation (3.46), (iii) is Fubini's theorem, and (iv) uses the fact that $\mathcal{P}$ is centered and isotropic, whereas $v^{(i)}$ are unit vectors. □

Finally, we establish the mutual information recursion from Theorem 4.3.

Proof of Theorem 4.3. The first equation of the theorem, Equation 4.16, follows directly from Proposition 4.1, Proposition 4.2, and the choice of loss function $\mathcal{L}(v^{(k+1)},\theta) = \mathbb{I}(\langle v^{(k+1)},\theta\rangle \geq \tau)$. By integrating Equation 4.16, we have for any τ^*

$$\begin{aligned}
\tau_{k+1} &\leq \tau^* + \int_{t=\tau^*}^{d}\left(\frac{\log 2 + \frac{\lambda^2}{2}(k+\sum_{i=1}^{k}\tau_i)}{C_1 t - C_2}\right)dt \\
&= \tau^* + \frac{\log 2 + \frac{\lambda^2}{2}(k+\sum_{i=1}^{k}\tau_i)}{C_1}\log\frac{d}{\tau^* - C_2/C_1}\,.
\end{aligned}$$

Setting $\tau^* = C_2/C_1 + \frac{\log 2+\frac{\lambda^2}{2}(k+\sum_{i=1}^{k}\tau_i)}{C_1}$ yields

$$\tau_{k+1} \leq \frac{C_2}{C_1} + \frac{\log 2 + \frac{\lambda^2}{2}(k+\sum_{i=1}^{k}\tau_i)}{C_1}\left(1 + \log\frac{C_1 d}{\log 2 + \frac{\lambda^2}{2}(k+\sum_{i=1}^{k}\tau_i)}\right).$$

□

C.2 General Upper Bound on Likelihood Second Moments

Again, fix $\tau_1,\ldots,\tau_{T+1}$, and define the event $A_\theta^k := \{\forall i\in[k],\ d\langle v^{(i)},\theta\rangle^2 \leq \tau_i\}$. We now proving the following lemma, which ecompasses both Propositions 5.2 and 6.2:

Lemma C.2 (Generic Upper Bound on Likelihood Ratios)**.** *Let $u,s\in\mathcal{S}^{d-1}$. Define the expected likelihood ratios:*

$$g_i(u,s;\{v^{(j)}\}_{1\leq j\leq i}) := \mathbb{E}_{\mathbf{P}_0}\left[\frac{\mathrm{d}\mathbf{P}_u(Z_i|Z_{i-1})\cdot\mathrm{d}\mathbf{P}_s(Z_i|Z_{i-1})}{\left(\mathrm{d}\mathbf{P}_0(Z_i|Z_{i-1})\right)^2}\Big|\{v^{(j)}\}_{1\leq j\leq i}\right]. \tag{3.47}$$

Then, letting $\mathcal{V}_\theta^k := \{v^{(1)}, \dots, v^{(k)} : \forall i \in [k],\ d\langle v^{(i)}, \theta\rangle^2 \le \tau_i\}$, *we have*

$$\mathbb{E}_{\mathbf{P}_0}\left[\frac{\mathrm{d}\mathbf{P}_u(Z_k; A_u^k)\cdot \mathrm{d}\mathbf{P}_s(Z_k; A_s^k)}{\left(\mathrm{d}\mathbf{P}_0(Z_k)\right)^2}\right] \quad \le \quad \sup_{v^{(1)},\dots,v^{(k)} \in \mathcal{V}_u^k \cap \mathcal{V}_s^k} \prod_{i=1}^k g_i(u, s; \{v^{(j)}\}_{1\le j\le i}). \tag{3.48}$$

Proof of Lemma C.2. Introduce the shorthand $v^{(1:i)} = \{v^{(j)}\}_{1\le j\le i}$, with the convention $v^{(1:0)} = \emptyset$. Fixing u and s, $A^j = A_u^j \cap A_s^j$ for $1 \le j \le k$, with the convention A^0 is the entire event space. We also define the "tail set"

$$\mathcal{V}_{j+1}^k(Z_{j-1}) := \{v^{(j+1)}, \dots, v^{(k)} : \forall i \in \{j+1, k\}, \theta \in \{u, s\} : d\langle v^{(i)}, \theta\rangle^2 \le \tau_i\}. \tag{3.49}$$

Finally, we define a partial supremum over the products of the terms $g_i(u, s; \{v^{(j)}\}_{1\le j\le i})$:

$$G_j(v^{(1:j)}) = \sup_{\widetilde{v}^{(j+1)},\dots,\widetilde{v}^{(k)} \in \mathcal{V}_{j+1}^k(Z_{j-1})} \prod_{i=j+1}^k g_i(u, s; \{v^{(1:j)}, \widetilde{v}^{(j+1:i)}\}). \tag{3.50}$$

adopting the convention $G_k(v^{(1:k)}) = 1$. We observe also that

$$G_0(v^{(1:0)}) = \sup_{\widetilde{v}^{(1)},\dots,\widetilde{v}^{(k)} \in \mathcal{V}_u^k \cap \mathcal{V}_s^k} \prod_{i=1}^k g_i(u, s; \widetilde{v}^{(1:k)}). \tag{3.51}$$

For any $j \in [k]$, we compute

$$\begin{aligned}
& \mathbb{E}_{\mathbf{P}_0}\left[G_j(v^{(1:j)})\frac{\mathrm{d}\mathbf{P}_u(Z_j)\cdot \mathrm{d}\mathbf{P}_s(Z_j)}{\left(d\mathbf{P}_0(Z_j)\right)^2}\mathbb{I}(A^j)\right] \\
= \quad & \mathbb{E}_{\mathbf{P}_0}\left[\mathbb{E}\left[G_j(v^{(1:j)})\frac{\mathrm{d}\mathbf{P}_u(Z_j)\cdot \mathrm{d}\mathbf{P}_s(Z_j)}{\left(\mathrm{d}\mathbf{P}_0(Z_j)\right)^2}\mathbb{I}(A^j)\big|Z_{j-1}\right]\right] \\
= \quad & \mathbb{E}_{\mathbf{P}_0}\left[\mathbb{E}\left[G_j(v^{(1:j)})\frac{\mathrm{d}\mathbf{P}_u(Z_{j-1})\cdot \mathrm{d}\mathbf{P}_s(Z_{j-1})}{\left(\mathrm{d}\mathbf{P}_0(Z_{j-1})\right)^2}\cdot\frac{\mathrm{d}\mathbf{P}_u(Z_j|Z_{j-1})\cdot \mathrm{d}\mathbf{P}_s(Z_j|Z_{j-1})}{\left(\mathrm{d}\mathbf{P}_0(Z_j|Z_{j-1})\right)^2}\mathbb{I}(A^j)\big|Z_{j-1}\right]\right] \\
\overset{(i)}{=} \quad & \mathbb{E}_{\mathbf{P}_0}\left[\frac{\mathrm{d}\mathbf{P}_u(Z_{j-1})\cdot \mathrm{d}\mathbf{P}_s(Z_{j-1})}{\left(\mathrm{d}\mathbf{P}_0(Z_{j-1})\right)^2}\cdot\mathbb{E}\left[G_j(v^{(1:j)})\frac{\mathrm{d}\mathbf{P}_u(Z_j|Z_{j-1})\cdot \mathrm{d}\mathbf{P}_s(Z_j|Z_{j-1})}{\left(\mathrm{d}\mathbf{P}_0(Z_j|Z_{j-1})\right)^2}\mathbb{I}(A^j)\big|Z_{j-1}\right]\right] \\
\overset{(ii)}{=} \quad & \mathbb{E}_{\mathbf{P}_0}\left[\frac{\mathrm{d}\mathbf{P}_u(Z_{j-1})\cdot \mathrm{d}\mathbf{P}_s(Z_{j-1})}{\left(\mathrm{d}\mathbf{P}_0(Z_{j-1})\right)^2}\cdot\mathbb{I}(A^j)G_j(v^{(1:j)})\mathbb{E}\left[\frac{\mathrm{d}\mathbf{P}_u(Z_j|Z_{j-1})\cdot \mathrm{d}\mathbf{P}_s(Z_j|Z_{j-1})}{\left(\mathrm{d}\mathbf{P}_0(Z_j|Z_{j-1})\right)^2}\big|Z_{j-1}\right]\right].
\end{aligned}$$

Where (i) follows since the densities $\frac{\mathrm{d}\mathbf{P}_u(Z_{j-1})\cdot \mathrm{d}\mathbf{P}_s(Z_{j-1})}{\left(\mathrm{d}\mathbf{P}_0(Z_{j-1})\right)^2}$ are functions of Z_{j-1}, and (ii) follows since A^j and G_j depend only on $v^{(1)}, \dots, v^{(j)}$, which is Z_{j-1} measurable because $\{v^{(1)}, w^{(1)}, \dots, v^{(j-1)}, w^{(j-1)}\}$ determine $v^{(j)}$. Finally, since Lemma 3.4 implies that the conditional law $\mathbf{P}_\theta[\cdot|Z_{j-1}]$ is independent of $w^{(1)}, \dots, w^{(j-1)}$ given $v^{(1)}, \dots, v^{(j-1)}$, we have

$$\begin{aligned}
& \mathbb{E}\left[\frac{\mathrm{d}\mathbf{P}_u(Z_j|Z_{j-1})\cdot \mathrm{d}\mathbf{P}_s(Z_j|Z_{j-1})}{\left(\mathrm{d}\mathbf{P}_0(Z_j|Z_{j-1})\right)^2}\big|Z_{j-1}\right] \\
& \qquad = \mathbb{E}\left[\frac{\mathrm{d}\mathbf{P}_u(Z_j|Z_{j-1})\cdot \mathrm{d}\mathbf{P}_s(Z_j|Z_{j-1})}{\left(\mathrm{d}\mathbf{P}_0(Z_j|Z_{j-1})\right)^2}\big|\{v^{(1)}, \dots, v^{(j)}\}\right] := g_j(u, s; v^{(1:j)}).
\end{aligned} \tag{3.52}$$

Moreover, on A^j, we have that $d\langle v^{(j)}, \theta\rangle^2 \le \tau_j$ for $\theta \in \{u, s\}$, which implies

$$\begin{aligned}
\mathbb{I}(A^j)G_j(v^{(1:j)})g_j(u,s;v^{(1:j)}) &= \mathbb{I}(A^j)g_j(u,s;v^{(1:j)}) \sup_{\widetilde{v}^{(j+1)},\dots,\widetilde{v}^{(k)}\in\mathcal{V}_{j+1}^k(Z_{j-1})} \prod_{i=j+1}^{k} g_i(u,s;\{v^{(1:j)},\widetilde{v}^{(j+1:i)}\}) \\
&\le \mathbb{I}(A^j) \sup_{\widetilde{v}^{(j)},\dots,\widetilde{v}^{(k)}\in\mathcal{V}_{j}^k(Z_{j-1})} \prod_{i=j}^{k} g_i(u,s;\{v^{(1:j-1)},\widetilde{v}^{(j+1:i)}\}) \\
&= \mathbb{I}(A^j)G_{j-1}(v^{(1:j-1)}) \\
&\le \mathbb{I}(A^{j-1})G_{j-1}(v^{(1:j-1)}).
\end{aligned}$$

where the last inequality follows since $A^j \subset A^{j-1}$. Altogether, we have proven

$$\mathbb{E}_{\mathbf{P}_0}\left[G_j(v^{(1:j)})\frac{\mathrm{d}\mathbf{P}_u(Z_j)\cdot\mathrm{d}\mathbf{P}_s(Z_j)}{(\mathrm{d}\mathbf{P}_0(Z_j))^2}\mathbb{I}(A^j)\right] \le \mathbb{E}_{\mathbf{P}_0}\left[G_{j-1}(v^{(1:j-1)})\frac{\mathrm{d}\mathbf{P}_u(Z_{j-1})\cdot\mathrm{d}\mathbf{P}_s(Z_{j-1})}{(\mathrm{d}\mathbf{P}_0(Z_{j-1}))^2}\mathbb{I}(A^{j-1})\right] \tag{3.53}$$

To conclude, we use the fact that $G_k(v^{(1:k)}) = 1$ and that $Z_0 = \emptyset$ to render

$$\begin{aligned}
\mathbb{E}_{\mathbf{P}_0}\left[\frac{\mathrm{d}\mathbf{P}_u(Z_k)\cdot\mathrm{d}\mathbf{P}_s(Z_k)}{(\mathrm{d}\mathbf{P}_0(Z_k))^2}\mathbb{I}(A^k)\right] &= \mathbb{E}_{\mathbf{P}_0}\left[G_k(v^{(1:k)})\frac{\mathrm{d}\mathbf{P}_u(Z_k)\cdot\mathrm{d}\mathbf{P}_s(Z_k)}{(\mathrm{d}\mathbf{P}_0(Z_k))^2}\mathbb{I}(A^k)\right] \\
&\le \mathbb{E}_{\mathbf{P}_0}\left[G_0(v^{(1:0)})\frac{\mathrm{d}\mathbf{P}_u(Z_0)\cdot\mathrm{d}\mathbf{P}_s(Z_0)}{(\mathrm{d}\mathbf{P}_0(Z_0))^2}\mathbb{I}(A^0)\right] \\
&= \sup_{\widetilde{v}^{(1)},\dots,\widetilde{v}^{(k)}\in\mathcal{V}_{j+1}^k(Z_{j-1})} \prod_{i=1}^{k} g_i(u,s;\widetilde{v}^{(1:k)}),
\end{aligned}$$

as needed.

□

C.3 Chi-Squared Computations

We begin with the following Lemma:

Lemma C.3. *The following identity holds,*

$$\mathbb{E}_{\mathbf{P}_0}\left[\frac{\mathrm{d}\mathbf{P}_u(Z_i|Z_{i-1})\mathrm{d}\mathbf{P}_s(Z_i|Z_{i-1})}{(\mathrm{d}\mathbf{P}_0(Z_i|Z_{i-1}))^2}\right] = \exp\left\{\lambda^2 d\langle v^{(i)},u\rangle\langle v^{(i)},s\rangle\left(u^\top P_{i-1}\Sigma_i^\dagger P_{i-1}s\right)\right\}. \tag{3.54}$$

Proof of Lemma C.3. We have that

$$\begin{aligned}
\frac{\mathrm{d}\mathbf{P}_u(Z_i|Z_{i-1})}{\mathrm{d}\mathbf{P}_0(Z_i|Z_{i-1})} &= \exp\left\{-\frac{1}{2}\left\|\lambda P_{i-1}uu^\top v^{(i)} - P_{i-1}w^{(i)}\right\|_{\Sigma_i/d}^2 + \frac{1}{2}\left\|P_{i-1}w^{(i)}\right\|_{\Sigma_i/d}^2\right\} \\
&= \exp\left\{-\frac{d\lambda^2}{2}\left\|P_{i-1}uu^\top v^{(i)}\right\|_{\Sigma_i}^2 + d\lambda u^\top v^{(i)}u^\top P_{i-1}\Sigma_i^\dagger P_{i-1}w^{(i)}\right\}.
\end{aligned}$$

Thus,

$$\begin{aligned}
\mathbb{E}_{\mathbf{P}_0}\left[\frac{\mathrm{d}\mathbf{P}_u(Z_i|Z_{i-1})\mathrm{d}\mathbf{P}_s(Z_i|Z_{i-1})}{(\mathrm{d}\mathbf{P}_0(Z_i|Z_{i-1}))^2}\right] &= e^{-\frac{d^2\lambda}{2}\left(\|P_{i-1}uu^\top v^{(i)}\|_{\Sigma_i}^2 + \|P_{i-1}ss^\top v^{(i)}\|_{\Sigma_i}^2\right)} \\
&\quad\times \mathbb{E}_{\mathbf{P}_0}\left[e^{d\lambda v^{(i)\top}(uu^\top+ss^\top)P_{i-1}\Sigma_i^\dagger P_{i-1}w^{(i)}}\right].
\end{aligned} \tag{3.55}$$

Now, note that if $Z \sim \mathcal{N}(0, I_d)$ is a standard gaussian vector, then we have that $w^{(i)} \stackrel{d}{=} \frac{1}{\sqrt{d}} \Sigma_i^{1/2} Z$, under $\mathbf{P}_0$. Thus,

$$\begin{aligned}
&\mathbb{E}_{\mathbf{P}_0}\left[e^{d\lambda v^{(i)\top}(uu^\top + ss^\top)P_{i-1}\Sigma_i^\dagger P_{i-1} w^{(i)}}\right] \\
&= \mathbb{E}_Z\left[e^{\sqrt{d}\lambda(\langle v^{(i)}, u\rangle u + \langle v^{(i)}, s\rangle s)^\top P_{i-1}\Sigma_i^{\dagger/2} P_{i-1} Z}\right] \\
&= e^{\frac{d\lambda^2}{2}\left\|P_{i-1}\Sigma_i^{\dagger/2}P_{i-1}(\langle v^{(i)}, u\rangle u + \langle v^{(i)}, s\rangle s)\right\|^2} \\
&\stackrel{(i)}{=} e^{\frac{d\lambda^2}{2}\left\|P_{i-1}(\langle v^{(i)}, u\rangle u + \langle v^{(i)}, s\rangle s)\right\|_{\Sigma_i}^2} \\
&= e^{\frac{d\lambda^2}{2}\left(\left\|\langle v^{(i)}, u\rangle P_{i-1}u\right\|_{\Sigma_i}^2 + \left\|\langle v^{(i)}, s\rangle P_{i-1}s\right\|_{\Sigma_i}^2\right) + d\lambda^2\langle v^{(i)}, u\rangle\langle v^{(i)}, s\rangle u^\top P_{i-1}\Sigma_i^\dagger P_{i-1}s},
\end{aligned} \tag{3.56}$$

where (i) uses the fact that P_{i-1} is symmetric, idempotent, and commutes with Σ_i. Combining Equations (3.55) and (3.56), we have

$$\mathbb{E}_{\mathbf{P}_0}\left[\frac{\mathrm{d}\mathbf{P}_u(Z_i|Z_{i-1})\mathrm{d}\mathbf{P}_s(Z_i|Z_{i-1})}{(\mathrm{d}\mathbf{P}_0(Z_i|Z_{i-1}))^2}\right] = e^{d\lambda^2\langle v^{(i)}, u\rangle\langle v^{(i)}, s\rangle u^\top P_{i-1}\Sigma_i^\dagger P_{i-1}s}.$$

□

Proposition C.4. *Fix $\tau_1, \ldots, \tau_T$, and let A_u^T denote the event that $d\langle v_i, u\rangle^2 \le \tau_i$ for all $i \in [T]$. Then,*

$$\mathbb{E}_{\mathbf{P}_0}\left[\frac{\mathrm{d}\mathbf{P}_u(Z_T; A_{u_1}^T)\mathrm{d}\mathbf{P}_s(Z_T; A_{u_2}^T)}{(\mathrm{d}\mathbf{P}_0(Z_T))^2}\right] \le e^{\lambda^2\{|\langle u_1, u_2\rangle|\sum_{i=1}^T \tau_i + \frac{1}{d}(\sum_{i=1}^T \tau_i)^2\}} \tag{3.57}$$

Proof of Proposition C.4. We have that

$$\begin{aligned}
(P_{i-1}\Sigma_i P_{i-1})^\dagger &= \left(P_{i-1}(I + v^{(i)}v^{(i)\top})P_{i-1}\right)^\dagger \\
&= P_{i-1}(I + v^{(i)}v^{(i)\top})^{-1}P_{i-1} \\
&= P_{i-1}(I - \frac{1}{2}v^{(i)}v^{(i)\top})P_{i-1} \\
&= P_{i-1} - \frac{1}{2}v^{(i)}v^{(i)\top} \\
&= I - \sum_{j=1}^{i-1} v^{(j)}v^{(j)\top} - \frac{1}{2}v^{(i)}v^{(i)\top},
\end{aligned}$$

where we use the fact that $P_{i-1}, v^{(i)}v^{(i)\top}$ and I commute throughout, and orthonormality of $v^{(1)}, \ldots, v^{(i)}$. Under $A_u^k \cap A_s^k$, we have that $|\langle v^{(i)}, \theta\rangle| \le \sqrt{\tau_i/d}$ for $\theta \in \{u, s\}$, which implies that

$$\begin{aligned}
&(v^{(i)\top}u)(v^{(i)\top}s)\{u^\top P_{i-1}\Sigma_i^\dagger P_{i-1}s\} \\
&= (v^{(i)\top}u)(v^{(i)\top}s)u^\top s - \frac{1}{2}(v^{(i)\top}u)^2(v^{(i)\top}s)^2 - (v^{(i)\top}u)(v^{(i)\top}s)\sum_{j=1}^{i-1}(v^{(j)\top}s)(v^{(j)\top}u) \\
&\le \frac{\tau_i}{d}|u^\top s| + \frac{1}{d^2}\sum_{j=1}^{i-1}\tau_i\tau_j.
\end{aligned}$$

Thus, by Lemma C.3, we have for any $(v^{(1)}, \dots, v^{(k)}) \in \mathcal{V}_u^k \cap \mathcal{V}_s^k$ that

$$g_i(u, s, 0; \{v^{(j)}\}_{1 \le j \le i}) \le e^{\lambda^2 \{\tau_i |u^\top s| + d^{-1} \sum_{j=1}^{i-1} \tau_i \tau_j\}}. \tag{3.58}$$

Equation (3.48) completes the demonstration, since for any $(v^{(1)}, \dots, v^{(k)}) \in \mathcal{V}_u^k \cap \mathcal{V}_s^k$,

$$\prod_{i=1}^{k} g_i(u, s, 0; \{v^{(j)}\}_{1 \le j \le i}) \le e^{\lambda^2 \sum_{i=1}^{k} \{\tau_i |u^\top s| + d^{-1} \sum_{j=1}^{i-1} \tau_i \tau_j\}} \le e^{\lambda^2 (\sum_{i=1}^{k} \tau_i |u^\top s| + d^{-1} (\sum_{i=1}^{k} \tau_i)^2)}. \tag{3.59}$$

□

D Additional Proofs

D.1 Proof of Theorem 5.3

For $k = 1$, we immediately have

$$\mathbb{E}_{\theta \sim \mathcal{P}} \mathbb{P}_\theta \left[d\langle v^{(1)}, \theta \rangle^2 > \tau_1 \right] \le \sup_{v \in \mathcal{S}^{d-1}} \mathbb{E}_{\theta \sim \mathcal{P}} \left[d\langle v, \theta \rangle^2 > \tau_1 \right]. \tag{4.60}$$

For $k \in \{2, \dots, T+1\}$, we define the action space $\mathcal{A} = \mathcal{S}^{d-1}$, and loss function $\mathcal{L}(v^{(k+1)}, \theta) = \mathbb{I}(\langle v^{(k+1)}, \theta \rangle \ge \tau_{k+1})$. Applying Proposition 5.1, then the inequality (5.26), we have

$$\begin{aligned}
&\mathbb{E}_{\theta \sim \mathcal{P}} \mathbb{P}_\theta \left[d\langle v^{(k+1)}, \theta \rangle^2 > \tau_k ; A_\theta^k \right] \\
&\le \sup_{u \in \mathcal{S}^{d-1}} \mathbb{P}_{\theta \sim \mathcal{P}} \left[d\langle u, \theta \rangle^2 \ge \tau_k \right] + \sqrt{\mathbb{E}_{\theta \sim \mathcal{P}} \mathbb{E}_{\mathbf{P}_0} \left[\left(\frac{\mathrm{d}\mathbf{P}_\theta(Z_k; A_\theta^k)}{\mathrm{d}\mathbf{P}_0(Z_k)} \right)^2 \right] \sup_{v \in \mathcal{S}^{d-1}} \mathbb{P}_{\theta \sim \mathcal{P}} \left[d\langle v, \theta \rangle^2 \ge \tau_k \right]} \\
&\le \sup_{u \in \mathcal{S}^{d-1}} \mathbb{P}_{\theta \sim \mathcal{P}} \left[d\langle u, \theta \rangle^2 \ge \tau_k \right] + e^{\frac{\lambda^2}{2} \sum_{i=1}^{k} \tau_i} \sqrt{\sup_{v \in \mathcal{S}^{d-1}} \mathbb{P}_{\theta \sim \mathcal{P}} \left[d\langle v, \theta \rangle^2 \ge \tau_k \right]} \\
&\le 2 e^{\frac{\lambda^2}{2} \sum_{i=1}^{k} \tau_i} \sqrt{\sup_{v \in \mathcal{S}^{d-1}} \mathbb{P}_{\theta \sim \mathcal{P}} \left[d\langle v, \theta \rangle^2 \ge \tau_k \right]}.
\end{aligned} \tag{4.61}$$

Hence,

$$\begin{aligned}
&\mathbb{E}_{\theta \sim \mathcal{P}} \mathbb{P}_\theta \left[\exists k \in [T+1] : \langle v^{(k)}, \theta \rangle^2 > \frac{\tau_k}{d} \right] \\
&= \sum_{k=1}^{T+1} \mathbb{E}_{\theta \sim \mathcal{P}} \mathbb{P}_\theta \left[\left\{ \langle v^{(k)}, \theta \rangle^2 > \frac{\tau_k}{d} \right\} \cap \left\{ \forall j < k, \langle v^{(j)}, \theta \rangle^2 \le \frac{\tau_j}{d} \right\} \right] \\
&= \sum_{k=1}^{T+1} \mathbb{E}_{\theta \sim \mathcal{P}} \mathbb{P}_\theta \left[\left\{ \langle v^{(k)}, \theta \rangle^2 > \frac{\tau_k}{d} \right\} ; A_\theta^{k-1} \right] \\
&\overset{(i)}{\le} \sup_{v \in \mathcal{S}^{d-1}} \mathbb{P}_{\theta \sim \mathcal{P}} [d\langle v, \theta \rangle^2 \ge \tau_1] + \sum_{k=1}^{T} \mathbb{E}_{\theta \sim \mathcal{P}} \mathbb{P}_\theta \left[\left\{ \langle v^{(k+1)}, \theta \rangle^2 > \frac{\tau_{k+1}}{d} \right\} ; A_\theta^{k-1} \right] \\
&\overset{(ii)}{\le} \sup_{v \in \mathcal{S}^{d-1}} \mathbb{P}_{\theta \sim \mathcal{P}} [d\langle v, \theta \rangle^2 \ge \tau_1] + 2 \sum_{k=1}^{T} e^{\frac{\lambda^2}{2} \sum_{i=1}^{k} \tau_i} \sqrt{\sup_{v \in \mathcal{S}^{d-1}} \mathbb{P}_{\theta \sim \mathcal{P}} \left[d\langle v, \theta \rangle^2 \ge \tau_{k+1} \right]}.
\end{aligned}$$

where (i) uses Equation (4.60), and (ii) uses (4.61). The theorem follows from a union bound by summing up the above display and combining with Equation (4.60).

D.2 Proof of Corollary 5.4

Set $\tau_1 = 2(\sqrt{\log(1/\delta)} + 1)^2$, so that

$$\frac{1}{2}(\sqrt{\tau_1} - \sqrt{2})^2 = \log(1/\delta).$$

Now, define $\tau_k \geq \sqrt{2}$ via

$$\frac{1}{2}(\sqrt{\tau_k} - \sqrt{2})^2 = \lambda^2 \sum_{i=1}^{k-1} \tau_i + (k-1)\tau_1.$$

Then, using the fact that $\mathbb{P}_{\theta\sim\mathcal{P}}[d\langle v, \theta\rangle^2 \geq \tau] \leq \exp(\frac{-1}{2}(\sqrt{\tau} - \sqrt{2})^2)$ from Lemma 4.4, Theorem 5.3 implies that

$$\begin{aligned}
& \mathbb{E}_{\theta\sim\mathcal{P}}\mathbb{P}_\theta\left[\exists k \in [T+1] : \langle v^{(k)}, \theta\rangle^2 > \frac{\tau_k}{d}\right] \\
\leq\ & \exp(\frac{-1}{2}(\sqrt{\tau_1} - \sqrt{2})) + \sum_{k=2}^{T+1} \sqrt{\exp(\lambda^2 \sum_{i=1}^{k-1} \tau_i - \frac{1}{2}(\sqrt{\tau_k} - \sqrt{2})^2)} \\
=\ & \exp(\frac{-1}{2}(\sqrt{\tau_1} - \sqrt{2})) + \sum_{k=2}^{T+1} \exp(-\frac{k-1}{2}\tau_1) \\
\leq\ & \exp(\frac{-1}{2}(\sqrt{\tau_1} - \sqrt{2})) + \sum_{k=2}^{T+1} \exp(-\frac{k-1}{2}(\tau_1 - \sqrt{2})^2) \\
=\ & \delta + \sum_{k=2}^{T+1} \delta^{k-1} \leq \frac{2\delta}{1-\delta}.
\end{aligned}$$

We now bound that rate at which our τ_k increase. Since τ_i are non-decreasing, we have

$$\frac{1}{2}(\sqrt{\tau_k} - \sqrt{2})^2 = \frac{\tau_k}{2}(1 - \sqrt{2/\tau_k}) \geq \frac{\tau_k}{2}(1 - \sqrt{2/\tau_1}).$$

And thus,

$$\tau_k \leq \frac{2}{1 - \sqrt{2/\tau_1}} \cdot (\lambda^2 \sum_{i=1}^{k-1} \tau_i + (k-1)\tau_1) \leq \frac{2(\lambda^2+1)}{1 - \sqrt{2/\tau_1}} \sum_{i=1}^{k-1} \tau_i.$$

For ease, set $\alpha = \frac{2(\lambda^2+1)}{1-\sqrt{2/\tau_1}}$, and consider the comparison sequence $\tau'_1 = \tau_1$, and $\tau'_k = \alpha \sum_{i=1}^{k-1} \tau'_i$. Then $\tau'_k \geq \tau_k$, and moreover, $\tau'_k \geq \alpha\tau'_{k-1}$, which implies that

$$\sum_{i=1}^{k-1} \tau'_i = \tau'_{k-1} \sum_{i=1}^{k-1} \frac{\tau'_i}{\tau'_{k-1}} \leq \tau_{k-1} \sum_{i=0}^{k-2} \alpha^{-i} \leq \frac{\tau'_{k-1}}{1 - 1/\alpha}.$$

Thus, $\tau_k \le \tau'_k = \alpha \sum_{i=1}^{k-1} \tau'_i \le \frac{\alpha}{1-1/\alpha}\tau'_{k-1} \le \tau'_1 \left(\frac{\alpha}{1-1/\alpha}\right)^{k-1} = \tau_1 \left(\frac{\alpha}{1-1/\alpha}\right)^{k-1}$. Finally, we bound

$$\begin{aligned}
\frac{\alpha}{1-1/\alpha} &= 2\lambda^2 \cdot \frac{1+1/\lambda^2}{(1-\frac{1-\sqrt{2/\tau_1}}{2(\lambda^2+1)})(1-\sqrt{2/\tau_1})} \\
&\le 2\lambda^2 \cdot \frac{1+1/\lambda^2}{(1-1/2\lambda^2)(1-\sqrt{1/(1+\sqrt{\log(1/\delta)})^2})} \\
&= 2\lambda^2 \cdot \frac{1+1/\lambda^2}{(1-1/2\lambda^2)(1-\sqrt{1/(1+\log(1/\delta))})} \\
&= 2\lambda^2 \cdot c(\lambda,\delta).
\end{aligned}$$

which implies that

$$\tau_k \le \tau_1 \left(2\lambda^2 \cdot c(\lambda,\delta)\right)^{k-1} = 2(\sqrt{\log(1/\delta)}+1)^2 \left(2\lambda^2 \cdot c(\lambda,\delta)\right)^{k-1}.$$

D.3 Proof of Proposition 3.1

Fix any $\lambda \ge 1$. Recall the function $c(\lambda,\delta) := (1+1/\lambda^2)\left\{(1-1/2\lambda^2)(1-1/\sqrt{1+\log(1/\delta)}\right\}^{-1}$ from Corollary 5.4. For $\delta \le 1/e$ and $\lambda \ge 1$, we have

$$c(\lambda,\delta) \le (1+1/\lambda^2)\left\{(1-1/2\lambda^2)(1-\sqrt{1+\log(1/\delta)}\right\}^{-1} \le \frac{2}{(1/2)(1-1/\sqrt{2})} := c'.$$

Next, fix any $\eta \ge 0$ and choose δ such that

$$\log(1/\delta) \le \left(\sqrt{\frac{d\eta}{2\left(2\lambda^2 \cdot c'\right)^T}} - 1\right)^2. \tag{4.62}$$

If we suppose that

$$\sqrt{\frac{d\eta}{2\left(2\lambda^2 \cdot c'\right)^T}} \ge 2. \tag{4.63}$$

This implies that $\delta \le 1/e$, so $c' \ge c(\lambda,\delta)$ and thus

$$\eta \ge \left(\left(2\lambda^2 \cdot c'\right)^T \cdot \frac{2\left(\sqrt{\log(1/\delta)}+1\right)^2}{d} \ge \left(\left(2\lambda^2 \cdot c(\lambda,\delta)\right)^T \cdot \frac{2\left(\sqrt{\log(1/\delta)}+1\right)^2}{d}. \right. \tag{4.64}$$

Then

$$\begin{aligned}
& \mathbb{E}_{\theta\sim\mathcal{P}}\mathbb{P}_\theta\left[\langle\widehat{v},\theta\rangle^2 \geq \eta\right] \\
\leq\ & \mathbb{E}_{\theta\sim\mathcal{P}}\mathbb{P}_\theta\left[\langle\widehat{v},\theta\rangle^2 \geq \left(\left(2\lambda^2\cdot c(\lambda,\delta)\right)^T\cdot\frac{2\left(\sqrt{\log(1/\delta)}+1\right)^2}{d}\right]\right. \\
\leq\ & \mathbb{E}_{\theta\sim\mathcal{P}}\mathbb{P}_\theta\left[\langle v^{(k)},\theta\rangle^2 \geq \left(\left(2\lambda^2\cdot c(\lambda,\delta)\right)^{k-1}\cdot\frac{2\left(\sqrt{\log(1/\delta)}+1\right)^2}{d}\forall k\in[T+1]\right]\right. \\
\overset{(i)}{\leq}\ & \frac{2\delta}{1-\delta}\overset{(ii)}{=}\frac{2\delta}{1-1/e},
\end{aligned}$$

where (i) uses Corollary 5.4 and (ii) uses that $\delta \leq 1/e$. Moreover, by Equation (4.62), we have

$$\begin{aligned}
\delta = \exp(-\log(1/\delta)) &\leq \exp\left(-\left(\sqrt{\frac{d\eta}{2\left(2\lambda^2\cdot c'\right)^T}}-1\right)^2\right) \\
&\leq \exp\left(-\left(\frac{1}{2}\sqrt{\frac{d\eta}{2\left(2\lambda^2\cdot c'\right)^T}}\right)^2\right) \\
&= \exp\left(-\frac{d\eta}{8\left(2\lambda^2\cdot c'\right)^T}\right).
\end{aligned}$$

Thus, if $\sqrt{\frac{d\eta}{2(2\lambda^2\cdot c')^T}} \geq 2$, we have

$$\mathbb{E}_{\theta\sim\mathcal{P}}\mathbb{P}_\theta\left[\langle\widehat{v},\theta\rangle^2 \geq \eta\right] \leq \frac{2}{1-e}\cdot\exp\left(-\frac{d\eta}{8\left(2\lambda^2\cdot c'\right)^T}\right).$$

On the other hand, if $\sqrt{\frac{d\eta}{2(2\lambda^2\cdot c')^T}} < 1$, then the right hand side of the above display is at least 1, so the result also holds vacuously.

D.4 Proof of Lemma 6.4

Let $c(\lambda,\delta)$ be as above, and fix $\delta \le 1/2$ and $\lambda \ge 1$. Then,

$$\begin{aligned}
\mathbb{E}_{\mathbf{P}_0}\left[\left(\frac{\mathrm{d}\overline{\mathbf{P}}[\cdot;\{A_\theta^T\}]}{\mathrm{d}\mathbf{P}_0}\right)^2\right] &= \mathbb{E}_{\theta,\theta'\sim\mathcal{P}}\mathbb{E}_{\mathbf{Q}}\left[\frac{\mathrm{d}\mathbf{P}_\theta[\cdot;A_\theta^T]\mathrm{d}\mathbf{P}_{\theta'}[\cdot;A_{\theta'}^T]}{(\mathrm{d}\mathbf{Q})^2}\right] \\
&= \mathbb{E}_{\theta,\theta'\sim\mathcal{P}} e^{\lambda^2\{|\langle\theta,\theta'\rangle|\sum_{i=1}^T\tau_i+\frac{1}{d}(\sum_{i=1}^T\tau_i)^2\}} \\
&= e^{\frac{\lambda^2}{d}(\sum_{i=1}^T\tau_i)^2}\cdot\mathbb{E}_{\theta,\theta'\sim\mathcal{P}} e^{\lambda^2(\sum_{i=1}^T\tau_i)\{|\langle\theta,\theta'\rangle|\}} \\
&\le e^{\frac{\lambda^2}{d}(\sum_{i=1}^T\tau_i)^2}\cdot e^{\frac{4\lambda^4}{d}(\sum_{i=1}^T\tau_i)^2+\lambda^2(\sum_{i=1}^T\tau_i)\sqrt{2/d}} \\
&= \exp\left\{\left(\frac{2\lambda^2\sqrt{1+1/4\lambda^2}}{\sqrt{d}}\left(\sum_{i=1}^T\tau_i\right)\right)^2+\frac{\sqrt{2}\lambda^2}{\sqrt{d}}\left(\sum_{i=1}^T\tau_i\right)\right\} \\
&\overset{(i)}{\le} \exp\left\{\frac{(2\sqrt{5/4}+\sqrt{2})\lambda^2}{\sqrt{d}}\cdot\sum_{i=1}^T\tau_i\right\} \le \exp\left\{\frac{\sqrt{7}\lambda^2}{\sqrt{d}}\cdot\sum_{i=1}^T\tau_i\right\},
\end{aligned}$$

where (i) holds as long as $\frac{(2\sqrt{5/4}+\sqrt{2})\lambda^2}{\sqrt{d}}\cdot\sum_{i=1}^T\tau_i \le 1$ and $\lambda \ge 1$. On the other hand, from Corollary 5.4, we can choose $\tau_1,\ldots,\tau_T$ such that, $\mathbb{E}_\theta\mathbf{P}_\theta[A_\theta^T] \ge \frac{2\delta}{1-\delta}$, and

$$\sum_{i=1}^T\tau_i \le 4(2\lambda^2\cdot c(\lambda,\delta))^{T-1}\left(\sqrt{\log(1/\delta)}+1\right)^2.$$

Then, as long as $4(2\lambda^2\cdot c(\lambda,\delta))^{T-1}\left(\sqrt{\log(1/\delta)}+1\right)^2 \le 1$, we have

$$\begin{aligned}
\mathbb{E}_{\mathbf{P}_0}\left[\left(\frac{\mathrm{d}\overline{\mathbf{P}}[\cdot;\{A_\theta^T\}]}{\mathrm{d}\mathbf{P}_0}\right)^2\right] &\le \exp\left\{\frac{4\sqrt{7}\lambda^2}{\sqrt{d}}(2\lambda^2\cdot c(\lambda,\delta))^{T-1}\left(\sqrt{\log(1/\delta)}+1\right)^2\right\} \\
&\overset{(i)}{\le} \exp\left\{\frac{4\sqrt{7}}{\sqrt{d}}(2\lambda^2\cdot c(\lambda,\delta))^{T}\left(\sqrt{\log(1/\delta)}+1\right)^2\right\}.
\end{aligned}$$

where (i) uses that $c(\lambda,\delta) \ge 1$. Since $\delta \le 1/2$ and $\lambda \ge 1$, we have that

$$c(\lambda,\delta) \le c'' := \max_{\delta\le 1/2,\lambda\ge 1} c(\lambda,\delta) = c(1,1/2) < \infty.$$

We can then bound

$$\mathbb{E}_{\mathbf{Q}}\left[\left(\frac{\mathrm{d}\overline{\mathbf{P}}[\cdot;\{A_\theta^T\}]}{\mathrm{d}\mathbf{Q}}\right)^2\right] \le \exp\left\{\frac{4\sqrt{7}\left(c''\lambda^2\right)^T}{\sqrt{d}}\cdot\left(\sqrt{\log(1/\delta)}+1\right)^2\right\}.$$

First, suppose that the quantity in the above exponential is less than $1/2$, then using the inequality $e^x - 1 \le 2x$ for $x \le 1/2$, we can bound

$$\mathbb{E}_{\mathbf{P}_0}\left[\left(\frac{\mathrm{d}\overline{\mathbf{P}}[\cdot;\{A_\theta^T\}]}{\mathrm{d}\mathbf{P}_0}\right)^2\right] \le \exp\left\{\frac{4\left(c''\lambda^2\right)^T}{\sqrt{d}}\cdot\left(\sqrt{\log(1/\delta)}+1\right)^2\right\} - 1 \le \frac{8\left(c''\lambda^2\right)^T}{\sqrt{d}}\cdot\left(\sqrt{\log(1/\delta)}+1\right)^2 \quad (4.65)$$

Take $\delta = \frac{\left(c''\lambda^2\right)^T}{\sqrt{d}}$. If $\delta \le 1/2$ and $\frac{8\left(c''\lambda^2\right)^T}{\sqrt{d}} \cdot \left(\sqrt{\log(1/\delta)}+1\right)^2 \le 1$, then

$$\begin{aligned}
\|\mathbf{P}_0 - \overline{\mathbf{P}}\|_{TV} &\overset{(i)}{\le} \frac{1}{2}\sqrt{\mathbb{E}_{\mathbf{Q}}\left[\left(\frac{\mathrm{d}\overline{\mathbf{P}}[\cdot;\{A_\theta^T\}]}{\mathrm{d}\mathbf{Q}}\right)^2\right] - 1} + \frac{\sqrt{2\mathbb{E}_{\theta\sim\mathcal{P}}\mathbf{P}_\theta[A_\theta^T]} + \mathbb{E}_{\theta\sim\mathcal{P}}\mathbf{P}_\theta[A_\theta^T]}{2} \\
&\overset{(ii)}{\le} \left(\sqrt{\log(1/\delta)}+1\right)\cdot\sqrt{\frac{2\left(\gamma''\lambda^2\right)^T}{\sqrt{d}}} + \frac{\sqrt{2\mathbb{E}_{\theta\sim\mathcal{P}}\mathbf{P}_\theta[A_\theta^T]} + \mathbb{E}_{\theta\sim\mathcal{P}}\mathbf{P}_\theta[A_\theta^T]}{2} \\
&\overset{(iii)}{\le} \left(\sqrt{\log(1/\delta)}+1\right)\cdot\sqrt{\frac{2\left(\gamma''\lambda^2\right)^T}{\sqrt{d}}} + \frac{\sqrt{2(\frac{2\delta}{1-\delta})} + \frac{2\delta}{1-\delta}}{2} \\
&\le \left(\sqrt{\log(1/\delta)}+1\right)\cdot\sqrt{\frac{2\left(\gamma''\lambda^2\right)^T}{\sqrt{d}}} + 4\sqrt{2\delta} \\
&\le \sqrt{2}(\log^{1/2}\frac{\sqrt{d}}{\left(\gamma''\lambda^2\right)^T} + 4)\cdot\sqrt{\frac{\left(\gamma''\lambda^2\right)^T}{\sqrt{d}}}.
\end{aligned}$$

where (i) uses Propostion 6.1, (ii) uses Equation (4.65), and (iii) uses the fact that $\mathbb{E}_{\theta\sim\mathcal{P}}\mathbf{P}_\theta[A_\theta^T] \le 2\delta/(1-\delta)$. On the other hand, if $\delta \ge 1/2$, or if $\frac{8\left(c''\lambda^2\right)^T}{\sqrt{d}} \cdot \left(\sqrt{\log(1/\delta)}+1\right)^2 \ge 1$, then the last line of the above display is at least 1, so the conclusion remained true because $\|\mathbf{P}_0 - \overline{\mathbf{P}}\|_{\mathrm{TV}} \le 1$.

E Information Theoretic Tools

The purpose of this section is to extend the theory of f-divergence based lower bounds ([30, 41, 54]) to handle non-normalized measures. In particular, this will allow us to prove analogues of the Fano-style Bayes Risk Lower Bounds in [26] which hold for truncated probability distributions, whose mass does not add up to one. The motivation for considering truncated measures is that we can restrict to parts of the probability space where the likelihood rations between alternatives do not "blow-up", allowing us to prove stronger bounds on the amount of information gained.

E.1 Preliminaries

Let μ be a finite, non-negative measure on a space $(\mathcal{X}, \mathcal{F})$, and denote $|\mu| = \mu(\mathcal{X})$. We say that μ is positive if $|\mu| > 0$. Abusive notation, we define $\mathbb{E}_{X\sim\mu}[f(X)]$ to denote integration of $f(X)$ with respect to the measure μ. Throughout, we will also let $f : (0,\infty) \to \mathbb{R}$ denote a convex function define $f'(\infty) := \lim_{x\to\infty} f(x)/x$. We allow the latter limit to be infinite.

Definition E.1. For a finite, non-negative measure μ and finite positive measure ν over the class $(\mathcal{X}, \mathcal{F})$, and a convex $f : (0,\infty) \to \mathbb{R}$, we define the (generalized) f-divergence between μ and ν as

$$D_f(\mu,\nu) := \int_{x\in\mathcal{X}:\mathrm{d}\nu(x)>0} f\left(\frac{\mathrm{d}\mu}{\mathrm{d}\nu}\right) d\nu + \mu\left(\{\mathrm{d}\nu = 0\}\right)\cdot f'(\infty) \tag{5.66}$$

with the convention $0 \cdot f'(\infty) = 0$.

Remark E.1. For a formal definition of how to interpret the notation $\mathrm{d}\mu, \mathrm{d}\nu$, see for example [47]. Recall that, for measures μ and ν on $(\mathcal{X}, \mathcal{F})$, we say $\mu \ll \nu$, or μ is absolutely continuous with respect to ν, if, for any set $A \in \mathcal{F}$, $\nu(A) = 0$ implies $\mu(A) = 0$. In this case, the Radon-Nikodym derivative $\mathrm{d}\mu/\mathrm{d}\nu$ is well defined [47], and $D_f(\mu, \nu) := \int_{x \in \mathcal{X}: \mathrm{d}\nu(x) > 0} f\left(\frac{\mathrm{d}\mu}{\mathrm{d}\nu}\right) \mathrm{d}\nu$.

The f-divergences play a substantial role in modern information theory, dating back to the work of Csiszár [30]. They also generalize many classical information divergence; for example, taking the function $f(x) = x^2 - 1$ yields the classical χ^2 divergence, and $f(x) = x \log x$ corresponds to the KL-divergence. Recall that $\chi^2 + 1$ divergence from Definition 5.1 corresponds to the case where $f(x) = x^2$, motivating the label "$\chi^2 + 1$-divergence". As defined in Definition 5.1, D_{χ^2+1} is not an f divergence in the traditional sense, since commonly one requires that $f(1) = 0$, and that $D_f(\cdot, \cdot)$ takes probability distributions as its arguments. In particular, this ensures that $0 = D_f(\mu, \mu) \leq \inf_\nu D_f(\mu, \nu)$ [54]. Luckily, the following lemma establishes that many of the nice properties of f divergences carry through when these restrictions are lifted.

Lemma E.1. *Let μ, ν be two finite positive measures on a space $(\mathcal{X}, \mathcal{F})$. Then*

1. ***Convexity:*** *$D_f(\mu, \nu)$ is jointly convex in μ and ν over the convex set $(\mu, \nu) : \mu \ll \nu$*
2. ***Distance-Like:*** *Suppose that $f'(\infty) \geq 0$. Then $D_f(\mu, \nu) \geq |\nu| f(|\mu|/|\nu|)$, which is attained when $\mathrm{d}\mu/\mathrm{d}\nu = |\mu|/|\nu|$.*
3. ***Normalization:*** *Define $f(x; p, q) = qf(\frac{p}{q}x)$. Then,*

$$D_f(\mu, \nu) = D_{f(x;|\mu|,|\nu|)}(\mu/|\mu|, \nu/|\nu|). \tag{5.67}$$

4. ***Linearity*** *$D_{\beta f + \alpha}(\mu, \nu) = \alpha|\nu| + \beta D_f(\mu, \nu)$*
5. ***Data-Processing:*** *Let Γ be a measurable map from $(\mathcal{X}, \mathcal{F})$ to $(\mathcal{Y}, \mathcal{G})$, and let $\mu\Gamma^{-1}$ denote the pullback measure on $\mathcal{Y}$ given by $\mu\Gamma^{-1}(B) = \mu(\Gamma^{-1}(B))$ for all $B \in \mathcal{G}$. Then,*

$$D_f(\mu, \nu) \geq D_f(\mu\Gamma^{-1}, \nu\Gamma^{-1}) \ .$$

It is often useful to consider f divergences on binary spaces (e.g. $|\mathcal{X}| = 2$). In this regime, lemma E.1 immediately implies the following corollary

Corollary E.2. *For $a \in [0, p]$, $b \in [0, q]$, let $\mathbf{Q}_{a,p}$ (resp. $\mathbf{Q}_{b,q}$) denote the measures on $\{0, 1\}$ which place mass a (resp. b) on 1, and $p - a$ (resp. $q - b$) on $\{0, 1\}$, and define*

$$\phi_f(a, b; p, q) = D_f(\mathbf{Q}_{a,p}, \mathbf{Q}_{b,q}). \tag{5.68}$$

Then, For $a \in [0, q], b \in [0, q]$,

$$\phi_f(a, b; p, q) = bf(\frac{a}{b}) + (q - b) \cdot f(\frac{p - a}{q - b}), \tag{5.69}$$

where the case $b = 0$ or $b = q$ is understood by taking the limits $b \to 0^+$ and $b \to q^-$. Moreover,

1. *Let μ, ν be measures on a space $(\mathcal{X}, \mathcal{F})$. For any event $A \in \mathcal{F}$, we have*

$$D_f(\mu, \nu) \geq \phi_f(\mu(A), \nu(A); |\mu|, |\nu|). \tag{5.70}$$

2. *The funciton $\phi_f(a,b;p,q)$ is jointly convex (and finite) in a and b for $(a,b) \in [0,p] \times (0,q)$, and*

$$\lim_{b\to 0} \phi_f(a,b;p,q) = \phi(a,0;p,q) \quad \lim_{b\to q} \phi_f(a,b;p,q) = \phi(a,q;p,q). \tag{5.71}$$

3. *As a function of a for fixed $b \in [0,q]$, $\phi_f(a,b;p,q)$ is minimized when $a = (p/q)b$, and is therefore nondecreasing for $(q/p)a \geq b$. As a function of b for fixed $a \in [0,p]$, $\phi_f(a,b;p,q)$ is minimized when $b = a(q/p)$ and is therefore nonincreasing for $(p/q)b \leq a$.*

Proof. The first point follows from applying the Data Processing Inequality from Part 5 of Lemma E.1 to the RHS of Equation (5.68). The second point is an analogue of Part 1 of Lemma E.1, but can be seen directly by noting that, if f is a convex function, the perspective map $(a,b) \mapsto bf(a/b)$ is convex (see [20]), and the limits follow from direct computation and the definition of $f'(\infty)$. The third point of Lemma E.1 to the RHS of Equation (5.68), and noting that 1-d convex functions are non-increasing to the left (resp. non-decreasing to the right) of their minimizers. □

E.2 A Generalized Bayes-Risk Lower Bound

With this in hand, we can prove generalization of the Bayes risk lower bounds from [26]. We consider a space of actions $\mathcal{A}$ (think queries $v^{(i)}$), and a $\{0,1\}$-loss function $\mathcal{L} : \mathcal{A} \times \Theta \to \{0,1\}$. As in [26], we consider an decision rule $\mathfrak{a} : \mathcal{X} \to \mathcal{A}$, and study $\mathcal{L}(\mathfrak{a}(X),\theta)$, where X is drawn from a measure μ_θ and θ comes from a prior $\mathcal{P}$ over a space $(\Omega,\mathcal{G})$. Rather than lower bounding the probability that $\mathcal{L}(\mathfrak{a}(X),\theta)$ is equal to one 1 (i.e., risk), we *upper bound* the probability that $\mathcal{L}((X),\theta)$ is equal 0, which we call the *value*. This will be easier to work with for our purposes, and makes more sense semantically when $|\mu_\theta| \leq 1$.

Theorem E.3 (Generalized Bayes Risk Lower Bound)**.** *Let $\mathcal{L} : \mathcal{A} \times \Theta \to \{0,1\}$, let $\mathfrak{a}$ denote a decision rule from $\mathcal{X} \to \mathcal{A}$, let $\mathcal{P}$ be a probability distribution over $(\Theta,\mathcal{G})$, let ν and $\{\mu_\theta\}$ be a family of finite measures over $(\mathcal{X},\mathcal{F})$ such that $\mathbb{E}_{\theta\sim\mathcal{P}}|\mu_\theta| > 0$. Define $p = \mathbb{E}_{\theta\sim\mathcal{P}}|\mu_\theta|$, $q = |\nu|$ and let $V^* = \sup_{\mathfrak{a}} \mathbb{E}_{\theta\sim\mathcal{P}}\mu_\theta[\{L(\mathfrak{a}(X),\theta) = 0\}]$, $V_0 = \sup_{a\in\mathcal{A}} \mathbb{P}_{\theta\sim\mathcal{P}}[L(a,\theta) = 0]$. Then,*

$$\textit{either} \quad V^* \leq pV_0 \quad \textit{or} \quad \mathbb{E}_{\theta\sim\mathcal{P}} D_f(\mu_\theta,\nu) \geq \phi_f(V_*, q\cdot V_0; p, q). \tag{5.72}$$

Proof. We follow along the lines of the proofs of Lemma 3 and Theorem 2 [26], but first we introduce some notation. Let $\mathcal{P}\otimes\nu$ denote the product measure between $\mathcal{P}$ and ν, and let $\mathcal{P} * \{\mu_\theta\}$ denote the coupled measure with density $(\mathcal{P} * \{\mu_\theta\})(\theta, \mathrm{X}) = \mathrm{d}\mathcal{P}(\theta)\cdot\mu_\theta(\mathrm{X})$. Also, given a measure η on $(\Theta,\mathcal{G})\times(\mathcal{X},\mathcal{F})$, define the *value* of $\mathfrak{a}$ as the mass of all pairs (θ,X) which incur zero loss

$$V_{\mathfrak{a}}(\eta) = \eta(\{(\theta,X) : \mathbb{I}(L(\mathfrak{a}(X),\theta) = 0\}). \tag{5.73}$$

Defining the event $A := \{(\theta,X) : \mathbb{I}(L(\mathfrak{a}(X),\theta) = 0\}$, we have

$$\begin{aligned}
\mathbb{E}_{\theta\sim\mathcal{P}} D_f(\mu_\theta,\nu) &= \int f\left(\frac{\mathrm{d}\mu_\theta}{\mathrm{d}\nu}\right) d(\mathcal{P}\otimes\nu) \\
&= \int f\left(\frac{\mathrm{d}\mathcal{P}\cdot\mathrm{d}\mu_\theta}{d\mathcal{P}\cdot\mathrm{d}\nu}\right) \mathrm{d}(\mathcal{P}\otimes\nu) \\
&= D_f\left(\mathcal{P} * \{\mu_\theta\}, (\mathcal{P}\otimes\nu)\right) \\
&\overset{(i)}{\geq} \phi_f\left((\mathcal{P} * \{\mu_\theta\})(A), \mathcal{P}\otimes\nu(A); |\mathcal{P} * \{\mu_\theta\}|, |\mathcal{P}\otimes\nu|\right) \\
&\overset{(ii)}{=} \phi_f\left(V_{\mathfrak{a}}(\mathcal{P} * \{\mu_\theta\}), V_{\mathfrak{a}}(\mathcal{P}\otimes\nu); p, q\right),
\end{aligned}$$

where (i) follows from the Data Processing inequality in Corollary E.2 Part 1, and for (ii) used the definition of $V^{\mathfrak{a}}$ and the definitions $p = |\mathcal{P} * \{\mu_\theta\}|$ and $q = |\mathcal{P} \otimes \nu| = |\mathcal{P}||\nu| = |\nu|$.

To wrap up, suppose that $V^* > pV_0$. We first note that $V_{\mathfrak{a}}(\mathcal{P} \otimes \nu) \leq |(\mathcal{P} \otimes \nu)| \cdot V_0 = qV_0$, since X and θ are independent under $\mathcal{P} \otimes \nu$. Moreover, for any $\epsilon > 0$, there exists a decision rule $\mathfrak{a}$ for which

$$p = |\mathcal{P} * \{\mu_\theta\}| \geq V_{\mathfrak{a}}(\mathcal{P} * \{\mu_\theta\}) > V^* - \epsilon. \tag{5.74}$$

Taking ϵ small enough $V^* - \epsilon > pV_0$, we have that

$$V_{\mathfrak{a}}(\mathcal{P} * \{\mu_\theta\}) > V^* - \epsilon > pV_0 = \frac{p}{q}(qV_0) \geq \frac{p}{q} V_{\mathfrak{a}}(\{\mathcal{P} \otimes \nu\}). \tag{5.75}$$

By Part 3 Corollary E.2, applied first to the 'b' argument and then to the 'a' argument, we have

$$\begin{aligned} \phi_f(V_{\mathfrak{a}}(\mathcal{P} * \{\mu_\theta\}); V_{\mathfrak{a}}(\{\mathcal{P} \otimes \nu\}); p, q) &\geq \phi_f(V_{\mathfrak{a}}(\mathcal{P} * \{\mu_\theta\}, qV_0; p, q) \\ &\geq \phi_f(V^* - \epsilon, qV_0; p, q). \end{aligned}$$

Since $\phi_f(a, b; p, q)$ is convex (Corollary E.2, Part 2), and therefore continuous, in its 'a' argument for $a \in [0, p]$, and since $V^* \leq p$, taking $\epsilon \to 0$ concludes. □

E.3 Application to Lower Bounds for Truncated Distributions

The flexibility to work with non-normalized measures allows us to prove lower bounds on the probability of taking a zero-loss action, restricted to some small parts of the space. We will also specialize Theorem E.2 to the case of the $\chi^2 + 1$ divergence defined in Definition 5.1.

Proposition E.4. *Let $\{\mathbf{P}_\theta\}_{\theta \in \Theta}$ denote a family of probability measures on a space $(\mathcal{X}, \mathcal{F})$ indexed by $\theta \in \Theta$. Let $\mathcal{P}$ denote a probability distribution on $(\Theta, \mathcal{G})$, and let $\{A_\theta\} \in \mathcal{F}$ such that $\{(\theta, A_\theta)\}$ is $(\mathcal{G}, \mathcal{F})$ measurable. Define $\mathbf{P}_\theta[\cdot; A_\theta]$ denote the subprobability measure given by*

$$\mathbf{P}_\theta[B; A_\theta] = \mathbf{P}_\theta[B \cap A_\theta]. \tag{5.76}$$

Given a loss function $\mathcal{L} : \mathcal{A} \times \Theta \to \{0, 1\}$, define

$$V^* = \sup_{\mathfrak{a}} \mathbb{E}_{\theta \sim \mathcal{P}} \mathbf{P}_\theta \left[\{L(\mathfrak{a}(X), \theta) = 0\} \cap A_\theta\right], \quad \text{and} \quad V_0 = \sup_{a \in \mathcal{A}} \mathbb{P}_{\theta \sim \mathcal{P}} \left[\{L(a, \theta) = 0\}\right]. \tag{5.77}$$

Letting $p = \mathbb{E}_{\theta \sim \mathcal{P}} \mathbf{P}_\theta[A_\theta]$, then either $V^* \geq pV_0$ or, for any probability measure $\mathbf{Q}$ on $(\mathcal{X}, \mathcal{F})$,

$$\mathbb{E}_{\theta \sim \mathcal{P}} D_f[\mathbf{P}_\theta[\cdot; A_\theta], \mathbf{Q}] \geq \phi_f(V^*, V_0; p, 1). \tag{5.78}$$

In particular, if we chose $f(x) = x^2$, then we have

$$V^* \leq V_0 + \sqrt{V_0(1 - V_0)\mathbb{E}_{\theta \sim \mathcal{P}} D_{\chi^2+1}[\mathbf{P}_\theta[\cdot; A_\theta], \mathbf{Q}]}. \tag{5.79}$$

Proof of Proposition 5.1. First we establish Equation (5.78), by applying Theorem E.3 with $\mu_\theta = \mathbb{P}_\theta[\cdot; A_\theta]$ and $\nu = \mathbf{Q}$. Since $\mathbf{Q}$ is a probability measure, $q = 1$. This gives

$$\mathbb{E}_{\theta \sim \mathcal{P}} D_f[\mathbb{P}_\theta[\cdot; A_\theta], \mathbf{Q}] \geq \phi_f(V^*, qV_0; p, q) = \phi_f(V^*, V_0; p, 1).$$

In the case that $f(x) = x^2$, we compute

$$\begin{aligned}
\phi_{\chi^2+1}(V^*, V_0; p, 1) &= \frac{(V^*)^2}{V_0} + \frac{(p-V^*)^2}{(1-V_0)} \\
&= \frac{(V^*)^2(1-V_0) + V_0(p-V^*)^2}{V_0(1-V_0)} \\
&= \frac{(V^*)^2 - 2pV_0V^* + p^2V_0}{V_0(1-V_0)} \\
&\overset{(i)}{\geq} \frac{(V^*)^2 - 2pV_0V^* + p^2V_0^2}{V_0(1-V_0)} \\
&= \frac{(V^* - pV_0)^2}{V_0(1-V_0)}.
\end{aligned}$$

where (i) uses that $V_0 \leq 1$. Thus, either $V^* \leq pV_0 \leq V_0$, or otherwise,

$$V^* \leq pV_0 + \sqrt{V_0(1-V_0)\mathbb{E}_{\theta\sim\mathcal{P}} D_f[\mathbf{P}_\theta[\cdot; A_\theta], \mathbf{Q}]} \tag{5.80}$$

$$\leq V_0 + \sqrt{V_0(1-V_0)\mathbb{E}_{\theta\sim\mathcal{P}} D_f[\mathbf{P}_\theta[\cdot; A_\theta], \mathbf{Q}]}\ . \tag{5.81}$$

□

Remark E.2. Examining the result of the above proposition, we see why it is so useful to consider the truncated measures $\mathbb{P}_\theta[\cdot; A_\theta]$ over their conditional analogues $\mathbb{P}_\theta[\cdot\big|A_\theta]$. This is because, if we instead proved a lower bound in terms of the latter, we would have to keep track of a normalization constants $\mathbb{P}_\theta[A_\theta]$, which would vary for all $\theta \in \Theta$. In particular, our case of interest considers the events $A_\theta^k = \{d\langle v^{(i)}, \theta\rangle^2 \leq \tau_i,\ \forall i \in [k]\}$ from Section 5. Because $v^{(i)}$ are chosen adaptively, it is quite difficult to control the normalization constant $\mathbb{P}_\theta[A_k^\theta]$ over all $\theta \in \mathcal{S}^{d-1}$.

E.4 Proof of Proposition 6.1

For ease of notation, let $\overline{\mathbf{P}}_A = \overline{\mathbf{P}}[\cdot; \{A_\theta\}]$. Note that $\overline{\mathbf{P}} - \overline{\mathbf{P}}_A \geq 0$, which impies that

$$\int |\mathrm{d}\overline{\mathbf{P}} - \mathrm{d}\overline{\mathbf{P}}_A| = \int \mathrm{d}\overline{\mathbf{P}} - \mathrm{d}\overline{\mathbf{P}}_A = 1 - \overline{\mathbf{P}}_A(\mathcal{X}) := 1 - p, \tag{5.82}$$

so by the triangle inequality

$$\begin{aligned}
&\|\mathbf{Q} - \overline{\mathbf{P}}\|_{TV} = \frac{1}{2}\int |\mathrm{d}\mathbf{Q}(x) - \mathrm{d}\overline{\mathbf{P}}(x)| \\
\leq\ &\frac{1}{2}\int |\mathrm{d}\mathbf{Q}(x) - \mathrm{d}\mathbf{P}_A(x)| + \frac{1}{2}\int |\mathrm{d}\overline{\mathbf{P}} - \mathrm{d}\overline{\mathbf{P}}_A| = \frac{1}{2}\int |\mathrm{d}\mathbf{Q}(x) - \mathrm{d}\mathbf{P}_A(x)| + \frac{1-p}{2},
\end{aligned}$$

Next, since $\mathbf{Q}$ is a probability measure,

$$\begin{aligned}\int |\mathrm{d}\mathbf{Q}(x) - \mathrm{d}\mathbf{P}_A(x)| &= \mathbb{E}_{\mathbf{Q}}|\frac{\mathrm{d}\overline{\mathbf{P}}_A}{\mathrm{d}\mathbf{Q}} - 1| \le \sqrt{\mathbb{E}_{\mathbf{Q}}|\frac{\mathrm{d}\overline{\mathbf{P}}_A}{\mathrm{d}\mathbf{Q}} - 1|^2} \\ &= \sqrt{\mathbb{E}_{\mathbf{Q}}|\frac{\mathrm{d}\overline{\mathbf{P}}_A}{\mathrm{d}\mathbf{Q}}|^2 + 1 - 2\overline{\mathbf{P}}_A(\mathcal{X})} = \sqrt{\mathbb{E}_{\mathbf{Q}}|\frac{d\overline{\mathbb{P}}_A}{\mathrm{d}\mathbb{Q}}|^2 + 1 - 2p} \\ &= \sqrt{\mathbb{E}_{\mathbf{Q}}|\frac{\mathrm{d}\overline{\mathbf{P}}_A}{\mathrm{d}\mathbf{Q}}|^2 - 1 + 2(1-p)} \le \sqrt{\mathbb{E}_{\mathbf{Q}}|\frac{\mathrm{d}\overline{\mathbf{P}}_A}{\mathrm{d}\mathbf{Q}}|^2 - 1} + \sqrt{2(1-p)}.\end{aligned}$$

Putting pieces together yields the proof.

E.5 Proof of Lemma E.1

The set $\{(\mu,\nu) : \mu \ll \nu\}$ is convex, since if $\alpha\nu_1(A) + (1-\alpha)\nu_2(A) = 0$, then $\nu_1(A) = \nu_2(A) = 0$, and thus if $\mu_1 \ll \nu_1$ nad $\mu_2 \ll \nu_2$, then $\alpha\mu_1(A) + (1-\alpha)\mu_2(A) = 0$. Moreover, the perspective map $(x,y) \to yf(x/y)$ is jointly for convex f [20], so that $\int f(\frac{\mathrm{d}\mu}{\mathrm{d}\nu})\mathrm{d}\nu$ is jointly convex in each argument.

For the second point , we see that that, by Jensen's inequality:

$$\begin{aligned}\int f(\frac{\mathrm{d}\mu}{\mathrm{d}\nu})\mathrm{d}\nu &= |\nu|\int f(\frac{\mathrm{d}\mu}{\mathrm{d}\nu})\frac{\mathrm{d}\nu}{|\nu|} \ge |\nu| f(\int \frac{\mathrm{d}\mu}{\mathrm{d}\nu}\frac{\mathrm{d}\nu}{|\nu|}) \\ &= |\nu| f(\frac{1}{|\nu|}\int \mathrm{d}\mu) = |\nu| f(\frac{|\mu|}{|\nu|}),\end{aligned}$$

so the result holds as long as $f'(\infty) \ge 0$.

Third, let $g(t) = f(t;p,q) = |\nu| f(t\frac{|\mu|}{|\nu|})$.Then $g'(\infty) = f'(\infty)\cdot|\mu|/|\nu|\cdot|\nu| = |\mu|f'(\infty)$. Thus,

$$\begin{aligned}D_f(\mu,\nu) &= \int f(\frac{\mathrm{d}\mu}{\mathrm{d}\nu})d\nu + \mu(\{\mathrm{d}\nu = 0\})f'(\infty) \\ &= \int |\nu| f(\frac{|\mu|}{|\nu|}\cdot\frac{\mathrm{d}(\mu/|\mu|)}{\mathrm{d}(\nu/|\nu|)})\cdot\mathrm{d}\nu/|\nu| + (\frac{\mu}{|\mu|})(\{d\nu = 0\})\cdot|\mu|f'(\infty) \\ &= \int g(\frac{\mathrm{d}(\mu/|\mu|)}{\mathrm{d}(\nu/|\nu|)})\mathrm{d}\nu/|\nu| + \frac{\mu}{|\mu|}(\{\mathrm{d}\nu = 0\})g'(\infty) \\ &= D_g(\mu/|\mu|, \nu/|\nu|),\end{aligned}$$

as needed. For the fourth point point, note that for any constant $\alpha, (f+\alpha)'(\infty) = f(\infty)$. Thus,

$$\begin{aligned}D_{f+\alpha}(\mu,\nu) &= \int\{f(\frac{\mathrm{d}\mu}{\mathrm{d}\nu}) + \alpha\}d\nu + \mu(\{\mathrm{d}\nu = 0\})(f+\alpha)'(\infty) \\ &= \alpha|\nu| + \int f(\frac{\mathrm{d}\mu}{\mathrm{d}\nu})\mathrm{d}\nu + \mu(\{\mathrm{d}\nu = 0\})(f)'(\infty) = \alpha|\nu| + D_f(\mu,\nu).\end{aligned}$$

Similarly, since $(\beta f)'(\infty) = \beta f'(\infty)$, one has

$$\begin{aligned}D_{\beta f}(\mu,\nu) &= \int\{\beta f(\frac{\mathrm{d}\mu}{\mathrm{d}\nu}) + \alpha\}\mathrm{d}\nu + \mu(\{\mathrm{d}\nu = 0\})(\beta f)'(\infty) \\ &= \beta\int f(\frac{\mathrm{d}\mu}{\mathrm{d}\nu})\mathrm{d}\nu + \beta\mu(\{\mathrm{d}\nu = 0\})(f)'(\infty) = \beta D_f(\mu,\nu).\end{aligned}$$

Finally, the fifth point follows from the standard data-processing inequality [26] in the case when f is convex and $f(1) = 0$ and μ, ν are both probability distributions. By the previous bound, the inequality can be extended to convex f where $f(1)$ is not necessarily zero, and normalized μ, ν, by noting that

$$\begin{aligned} D_f(\mu,\nu) &= D_{f-f(1)}(\mu,\nu) + f(1)|\nu| \\ &\geq D_{f-f(1)}(\mu\Gamma^{-1},\nu\Gamma^{-1}) + f(1)|\nu| \quad \text{(classical data processing, e.g. Theorem 3.1 in Liese [54])} \\ &= D_{f-f(1)}(\mu\Gamma^{-1},\nu\Gamma^{-1}) + f(1)|\nu\Gamma^{-1}| \quad (\Gamma \text{ preserves total mass}) \\ &= D_f(\mu\Gamma^{-1},\nu\Gamma^{-1}). \end{aligned}$$

To generalize to arbitrary finite, positive measures, we note that the function $f(t;|\mu|,\nu)$ is convex, so

$$\begin{aligned} D_f(\mu,\nu) &= D_{f(;|\mu|,\nu)}(\mu/|\mu|,\nu/|\nu|) \\ &= D_{f(;|\mu|,\nu)}(\frac{\mu}{|\mu|}\Gamma^{-1},\frac{\nu}{|\nu|}\Gamma^{-1}) \\ &= D_f(\mu\Gamma^{-1},\nu\Gamma^{-1}). \end{aligned}$$

F Proof of Lemma 4.4

We begin by invoking a result from spherical isoperimetry:

Theorem F.1 (Spherical Isoperimetry, page 211 in [17])**.** *Let f be an 1-Lipschitz function on the sphere $\mathcal{S}^{d-1}$. Then,*

$$\mathbb{P}_{\theta\sim\mathcal{S}^{d-1}}\left[f(\theta) \geq \text{Median}(f) + t\right] \leq \sup_{A\subset\mathcal{S}^{d-1}} \mathbb{P}_{\theta\sim\mathcal{S}^{d-1}}\left[\theta \in A_t\right] \leq e^{-dt^2/2}. \tag{6.83}$$

where $A_t := \{\theta : \exists\ \theta' \in A \text{ such that } \|\theta - \theta'\|_2 \leq t\}$.

We can now prove Lemma 4.4. Observe that $\text{Median}(|\langle\theta, v\rangle|) \leq \sqrt{2/d}$, since by Markov's inequality,

$$\mathbb{P}[|\langle\theta, v\rangle| \geq \sqrt{2/d}] \leq \mathbb{E}[|\langle\theta, v\rangle|^2]/(2/d) = d/2 \cdot (1/d) = 1/2. \tag{6.84}$$

Since $\theta \mapsto |\langle\theta, v\rangle|$ is 1-Lipschitz, Theorem F.1 yields

$$\mathbb{P}[|\langle\theta, v\rangle| \geq \sqrt{2/d} + t] \leq e^{-dt^2/2}. \tag{6.85}$$

Replacing t with $\sqrt{t/d}$ concludes the proof.